AN ARGUMENT FOR UNIQUENESS ABOUT EVIDENTIAL SUPPORT

Sinan Dogramaci
The University of Texas at Austin

Sophie Horowitz
University of Massachusetts, Amherst

Abstract

White, Christensen, and Feldman have recently endorsed uniqueness, the thesis that given the same total evidence, two rational subjects cannot hold different views. Kelly, Schoenfield, and Meacham argue that White and others have at best only supported the weaker, merely intrapersonal view that, given the total evidence, there are no two views which a *single* rational agent could take. Here, we give a new argument for uniqueness, an argument with deliberate focus on the interpersonal element of the thesis. Our argument is that the best explanation of the value of promoting rationality is an explanation that entails uniqueness.

Introduction

We give an argument for **uniqueness**: given the same total evidence, two subjects cannot rationally hold different views. We intend "view" to cover both partial and full beliefs.

Many authors support uniqueness's negation, **permissivism**: given the same total evidence, two rational subjects can hold different views. Supporters of permissivism (call them "permissivists") include Rosen (2001), Schoenfield (2014), Kelly (2010 & 2013), Meacham (2014), and Ballantyne & Coffman (2011).

Authors sympathetic to uniqueness (call them "impermissivists") include Feldman (2007), who presents some cases offering intuitive support for uniqueness, Christensen (2007), who operates with uniqueness as an "attractive" "assumption" (p.211), and White (2005 & 2013), who

constructs several arguments aimed at exposing absurdities that follow from permissivism.

Kelly (2013) has argued that White's arguments fail to apply to uniqueness, but rather only support **uniqueness***, the thesis that, given the total evidence, there are no two views either of which a *single* rational agent could take. In other words, Kelly claims White may have supported an *intra*personal uniqueness thesis, but not the original, more ambitious and interesting *inter*personal thesis of uniqueness. Many of Schoenfield's (2014) and Meacham's (2014) objections to White also support this diagnosis.

Here, we offer a new positive argument for uniqueness, with a focus on the interpersonal character of the thesis. The form of our argument is an *inference to the best explanation*. Let's begin, then, by introducing the datum we propose to best explain.

The Datum

The datum we'll explain, i.e. the premise of our overall argument, is this: *our social practice of epistemically evaluating one another's beliefs has value*. Let's now clarify and elaborate that.

As a matter of contingent, actual fact, we use epistemic evaluations to discourage irrational beliefs and encourage rational beliefs. We discourage certain beliefs by using criticisms like "That's irrational!", or more often we use more colloquial language like "You don't know what you're talking about!" or "Don't be so stupid!", where context makes clear the evaluation is distinctively epistemic. Third-personal evaluations may be more common, and more diplomatic and, thus, more effective at influencing audiences than second-personal evaluations; you might have better luck influencing your audience by criticizing a third-party, maybe saying something like "Aren't these climate change deniers an embarrassment?" Positive evaluations, correspondingly, serve to encourage certain beliefs. And, all this epistemic criticism and praise can be applied to degrees of belief, e.g. when we praise someone's low confidence, maybe with something like, "Pat has a healthy skepticism about new age herbal remedies".

In a typical epistemic evaluation, an evaluative predicate is applied to a particular belief, maybe as held by a particular person in a particular context. However, it seems clear on reflection, and it is agreed by many epistemologists, that our ordinary epistemic evaluations of particular beliefs are true or false depending on the way the belief was formed on the basis of the believer's evidence. Thus, the fundamental objects of epistemic evaluations are rules (or "processes", or "methods") that license certain beliefs given certain evidence. More abstractly, we might think of them as functions from total bodies of evidence to total doxastic states (Goldman (1986, ch.4), Pollock & Cruz (1999, p.123), Boghossian (2007)). Our datum,

then, is that there is some value in our social practice of promoting the use of certain epistemic rules, though we often go about this by only explicitly commenting on specific people's beliefs held in specific contexts.

The datum we hope to explain concerns a highly social practice, an interpersonal practice. It isn't just valuable to you that *you* are rational; it's valuable if *others* are rational as well. We see this as part of the job of epistemology: the right theory of rationality should explain not only what rationality is, but why it is valuable for us, and—our concern here—why it is valuable to promote rationality throughout one's community.

Exactly what kind of *value* do we take our epistemically evaluative practice to have? We believe our evaluative practice has a distinctly practical kind of value for us. Specifically, we take our evaluative practice to have instrumental value relative to our actual goals. (The practice may be valuable in other ways, as well, but we will focus on this one.) Our datum is that, in promoting epistemic rationality, we serve some goal of ours, some goal that we actually care to promote. We seek a non-trivial explanation of this datum, that is, we intend to show that the promotion of epistemic rationality serves, not simply the trivially served goal of promoting epistemic rationality, but some independent goal of ours.

Our view is that the practice of promoting epistemic rationality has instrumental value in the *actual* world, given our *actual* goals. This view is the datum of the inference-to-the-best-explanation argument we'll make, but, though it is thus a "datum", it is not an arbitrary and unsupported assumption for us. This view seems to us defeasibly reasonable on its face; the view could be proved wrong, but if it ever is, we'll lose our motivation to continue promoting rationality, and until then, we should defeasibly hold that our practices serve some desired end, even if ordinary people do not explicitly know how promoting rationality promotes a goal we independently care to serve. The alternative position that we've been serving no worthwhile end seems implausibly uncharitable to ourselves. We concede that possible strange cases can be concocted where promoting rationality only frustrates our goals, even our actual goals; but we do not claim, and we see no reason to claim, that promoting rationality *necessarily* serves our goals. Our claim is that the overall practice, on the whole, actually serves our actual goals. We mean to offer an explanation of how our practice does this, indeed of how our practice does this very *efficiently*. (In the next section, we'll add some criticisms of that alternative view that our evaluative practice simply lacks non-trivial instrumental value.)

A brief note for those familiar with some related literature: since we are explaining the instrumental value of our practice of *promoting* rationality, and not simply of *being* rational, our project is in no conflict with the critiques of the instrumental value of being rational given by Kelly (2003) and Berker (2013). Our position is also consistent with that of Madison (forthcoming), who argues that epistemic rationality, and presumably thus its promotion

too, has a non-instrumental value; there may be both kinds of value. Our project is in the spirit of Craig (1990), who asks what practical function is served by attributions of knowledge; however, our interest in this paper is in the function served by epistemic evaluations of rationality (or justification, warrant, reasonableness—these mean the same to us).

Three Inadequate Views of the Instrumental Value of Promoting Rationality

Shortly, we will offer an explanation that entails uniqueness. To support the inference to this explanation, and thereby to uniqueness, we must argue that this is the *best* explanation of how promoting rationality is instrumentally valuable, how it serves our goals. In order to help motivate and illustrate the comparative strength of our explanation, we lead up to it by considering some competing views of what, if anything, explains the instrumental value of promoting rationality, and showing why they are inadequate.

Here is a first candidate explanation: *promoting rationality serves our goals because, in doing so, we promote consistency among our interlocutors' beliefs.*

A first problem with this candidate is that it is, at best, very far from being a complete explanation of the value of our actual evaluative practices. It leaves us wondering why we criticize subjects whose beliefs are perfectly consistent. We actually do criticize subjects—whether actual people or hypothetical subjects of thought experiments—for making more substantive rational errors. For example, we criticize conspiracy theorists on rational grounds, even if their stories are consistent down to the last detail. We use epistemic evaluations to endorse taking appearances and testimony at face value, and reasoning according to certain statistical and inductive rules, including the projection of green rather than grue. And it is plausible that, in promoting standards of rationality that go beyond mere consistency, we must be serving some goal we actually have.

A second problem with this candidate explanation is that the promotion of consistency is not, plausibly, itself a goal that we have—nor does it promote other goals that we have. If we demand consistency of everyone, even skeptics, conspiracy theorists, and gruesome reasoners, it is not clear what goal of ours is served by this demand. So, the suggestion that we promote rationality because we care about promoting consistency does not adequately explain the value of our practice.

Next, let's consider the view that denies that the value of our evaluative practice can be given any non-trivial explanation: *the promotion of rationality serves no other goal than the trivial goal of making people more rational.*

Many epistemologists are attracted to this kind of answer; they are happy to see the value of rationality, if there is any at all, as *sui generis*; they see the value of rationality as a brute property, unexplainable as

instrumental value in the service of any other goal. It may be impossible for us to say anything that positively refutes this view without begging the question against those who are committed to it. But for those on the fence, we can say why we do not find this position plausible.

One thing to say is that, as a methodological point, this kind of answer should be our last resort. We are considering a universal human linguistic practice with a deep biological and cultural history: it is highly unlikely that it developed for no practical purpose. So it is worth looking at what explanations of this practical purpose we can find, and considering their plausibility, before positing brute facts or unexplainable kinds of value, or saying there is no value. At this point in the dialectic, it's plausible that promoting rationality *does* serve some non-trivial goal. And, as we'll shortly see, there *are* plausible candidate explanations for what this non-trivial instrumental value is. Rejecting the possibility of an explanation seems to us, in this case, premature.

We can also bring out the implausibility of this objection by comparing it with the corresponding view about the instrumental value of having and promoting knowledge. While it may seem that knowledge is valuable for its own sake, and while it may even seem to some that this is the *entire* explanation of the value of promoting it, there exists a far more substantive explanation of the instrumental value of not only having knowledge, but of promoting knowledge in others. The substantive explanation, in rough outline, runs as follows. Knowledge entails true belief, and true belief is instrumentally valuable, both in oneself and in others. (We always mean true belief *about matters we're actually interested in*.) True belief is instrumentally valuable in others, because so much of our information is derived from the accurate testimony of others; and true belief is instrumentally valuable in oneself, because actions that are based on true beliefs tend to serve our goals. (Functionalists like Lewis (1974, p.337) and Stalnaker (1984, p.15) suggest this link between actions, beliefs and goals/desires is necessary; here we only assume it is actual. See Horwich (1990/98, section 3.11) and Goldman (1992, p.164) for more endorsements.)

We can extend this truth-oriented explanation (of the instrumental value of knowledge and its promotion) even further. Plausibly, knowledge is not only true, but *reliably* true, perhaps on some "safety" model of reliability, including Williamson's (2000 & 2009, p.329) *non-reductive* notion of safety (that is to say, we do *not* assume the notion of safety can be fully explicated independently of our notion of knowledge). A belief's safe truth consists in its truth throughout some range of nearby possible worlds where that belief is held. (Exactly what's in the range is hard to say, hence the appeal of the non-reductive notion.) What is the instrumental value of this further feature of knowledge, beyond its mere factivity? The value is that, even when we are uncertain which world, and which imminent future, we actually find ourselves facing, if our beliefs are safely true then they will be actually true, no matter which nearby world turns out to be actual. Reliable, or safe, truth

is thus tremendously useful, in ourselves and in others we might ever need to rely on (Craig (1990, pp.19–20)).

Thus, the instrumental value of having and promoting knowledge can be substantively explained, in terms of the independent instrumental value of truth and reliability. We need not settle for saying (nothing more than) that knowledge is valuable for its own sake. We thus likewise doubt that we should settle for saying (nothing more than) that rational belief is valuable for its own sake, or not valuable at all.

The above discussion of the value of knowledge given in terms of truth and reliability may lead one to suspect that a parallel explanation can be given of the instrumental value of rationality. Thus, we have the third and last view we'll criticize here: *it is instrumentally valuable to promote rationality because, in doing so, we make it more likely that one another's beliefs will be true.* Or alternatively: *it is instrumentally valuable to promote rationality because, in doing so, we make one another's beliefs more reliable.*

We think this view is on the right track. We already observed that true beliefs and reliable beliefs are instrumentally valuable and worth promoting. And, in fact, we do think that using epistemic evaluations to promote rationality is instrumentally valuable because it promotes true belief, and we have true belief as a goal. Our criticism of this view is that, as it stands, it is missing a large core of the real explanation of *how* promoting rationality serves our truth goal. If this view, stated just as succinctly as in the italicized sentences in the last paragraph, gave the whole story, we would be left with an unexplained mystery as to why so many of us are so ready to criticize (as epistemically irrational) various characters who are known to be highly and robustly reliable, though unwittingly so, famous characters such as BonJour's (1980) clairvoyant, Lehrer's (1990) Mr. Truetemp, or Plantinga's (1993) brain lesion sufferer. As these cases show, our epistemically evaluative practice does not consider reliability a sufficient condition for rationality. These characters' would have highly reliable beliefs if they were allowed free rein to continue trusting their belief-forming methods, so the present candidate explanation cannot account for why we promote standards of rationality that exclude these characters. Is criticism of these characters misplaced or pointless? Is this a glitch in our practice, a kind of evaluation that serves no purpose? If some explanation accounts for the apparent value of promoting a standard of rationality that excludes these characters, that explanation is a better one. Is there any better explanation of the instrumental value of promoting our ordinary standard of rationality?

Our Proposed Explanation of the Value of Promoting Rationality

Something must explain why using epistemic evaluations to promote rationality is a worthwhile practice. It naturally seems that the explanation

should have something to do with a connection between rationality and truth or reliability. However, as we noted above, examples like BonJour's (1980) unwittingly reliable clairvoyant show that the belief-forming rules that yield rational beliefs cannot simply be identified with the reliable belief-forming rules.

The existence of examples like BonJour's clairvoyant now reveals that our original datum cries out for an explanation far more loudly than it initially appeared. There is something distinctly puzzling in our ordinary epistemically evaluative practice. The puzzle is this. Naturally, we want true belief, in ourselves and others, and we want to use ways of forming beliefs that are reliable, and so, naturally, we thus use "true" and "reliable" to correct each others' beliefs and modify each others' ways of forming beliefs. But, if our terms of epistemic evaluation, "[ir]rational" and various colloquial synonyms and antonyms, *differ in extension* from "[un]reliable", then *why* do we go around making epistemic evaluations of one another's beliefs, or ways of forming beliefs? Why do we use these epistemically evaluative terms to reinforce certain ways of forming beliefs, and to discourage or suppress other ways of forming beliefs, if we recognize that these ways *do not line up with the reliable ways*? It looks mysterious. A major virtue of the explanation we will propose of the practical, instrumental value of our epistemically evaluative practice is that it will solve this mystery.

To begin with, we make an observation: what the well-known examples like the clairvoyant still leave open, and what is anyway plausible, is that the rational rules are a *subset* of the rules that are actually reliable. This observation allows us to give the following explanation of the value of rationality.

The core idea of our offered explanation is this: *given uniqueness, the promotion of rationality is an efficient means of ensuring the reliability of testimony*. Our explanation thus will make the promotion of rationality an instrumentally valuable practice, given our goal of getting the truth (about matters of interest). Let us now elaborate our proposed explanation and its theoretical virtues. In particular, we will argue that our view allows us to treat one another as *epistemic surrogates*, and that it allows for an efficient *division of epistemic labor*.

First Virtue: Epistemic Surrogates

If uniqueness is true, then the promotion of rational belief-forming rules is the promotion of a fixed stock of belief-forming rules that are reliable and shared by rational reasoners. When rational reasoners thus coordinate upon the same reliable belief-forming rules, this has the result that rational reasoners are able to serve as each other's *epistemic surrogates*: the views that a rational testifier reports are the same reliable views that a rational audience

would have arrived at on their own, had the audience acquired the testifier's evidence and drawn their own conclusions.

Treating one another as epistemic surrogates is both an *effective* and *efficient* way of promoting our goal of getting to the truth (about matters of interest). It is *effective* because, given that one's surrogates share one's own (rational) epistemic rules or methods, their testimony will be reliable. (How do we *know* that they are reliable? In the same way we know that our own epistemic rules or methods are reliable. We are not skeptics; we assume that such a thing can be known. In fact, our view takes this basic anti-skeptical commitment and takes it a step farther: while any non-skeptic should admit that a rational agents can know that *her own* rules or methods are reliable, on our view she can also know that the *rational* rules are reliable.)

Of course, sometimes things are a bit more complicated, and we cannot simply take surrogates' testimony on board. A common type of case like this is when you offer me testimony about something I already have excellent evidence about—evidence that's much better than yours. Suppose you tell me that Bob is in his office, but I have just come back from looking in his office and it is empty. Even if you are an excellent surrogate, I should not trust your testimony in this case. Somewhat more complicated issues arise when our different evidence is equally good, but we come to opposite conclusions. And there are genuine puzzles regarding cases where we have the *same* evidence, but come to opposite conclusions, as the vast literature on peer disagreement brings out. (Given uniqueness, such cases arise only if one of us has made a rational mistake.) However, these more complex cases are not the norm. Most of the time, testimony is a way of learning about things that we have *not* investigated for ourselves—such as when you go down the hall by yourself, peek into Bob's doorway, and report back to me that he's in his office. And even in the more complex cases, hearing a surrogate's testimony will be a useful part of our overall efforts to converge onto the truth. (See Dogramaci (2012, fn 17) for further discussion of this issue. Greco & Hedden (forthcoming) discuss a view that is similar to Dogramaci's; they sidestep this issue by specifying that rationality ascriptions only commit us to deference in cases where we do not have relevant evidence that our testifier lacks.)

Promoting epistemic surrogates is also an *efficient* means of ensuring the reliability of testimony. If everyone conforms to a shared stock of epistemic rules, there is no need to investigate the reliability of particular testifiers. This is more efficient than the alternative: if people don't conform, we need to take the extra step of investigating or else risk gullibly accepting falsehoods. For example, consider a non-conformist who is presently offering us testimony that P. She might be reliable, like BonJour's clairvoyant, but she also might be quite unreliable. Her testimony cannot safely be taken at face value; to know whether she can be trusted, we must look into the track record of her belief-forming rules. Such an investigation costs resources, and by the time

we have concluded the investigation, we may well have found out whether P on our own. So responding to the testimony of non-conformists is very inefficient, even if it leads us to the truth in the end. Promoting conformity, which enables epistemic surrogacy, is far more efficient.

(But wait: isn't there need for another costly investigation—namely, an investigation into who is and is not a conformist—before we can take one another's testimony on board? Maybe so. But *promoting* conformity brings us toward a state where more and more people are reliable surrogates; so, our practice reduces the need for this investigation. As we succeed in promoting rationality, and thereby conformity, we make testimony safer and safer to trust. This is not so for views that deny uniqueness, where promoting rationality does not involve conformity. On those views, even if we were all to become rational, we would *still* need to check one another's reliability before trusting one another's testimony.)

Promoting, in this way, epistemic surrogates is thus the efficient way of ensuring the reliability of testimony. We can safely place as much trust in our surrogates' testimony as we do our own faculties, with no extra efforts or costs of vetting our testifiers' track-records.

Second Virtue: Division of Epistemic Labor

Having such trustworthy surrogates is extremely valuable, because if we can safely trust others' testimony on particular occasions, then we can cooperate with others for more efficient inquiry. We all want to acquire true beliefs, and no false beliefs, about various topics of interest. This task is efficiently pursued by a team of epistemic surrogates, one that functions like a team of parallel processors. Working together allows us to divide the labor of *collecting evidence* and the labor of *reasoning*.

Collecting evidence: On this parallel-processor model, different inquirers can go out and collect different batches of evidence, and the testimony they report back on the basis of this evidence can be safely trusted by the other surrogates. In the simplest case, a surrogate's testimony becomes as good as one's own eyes and ears, but without the effort of seeing and hearing for oneself. And even in more complex cases, as when two testifiers disagree, having epistemic surrogates is more efficient than the alternative. If we conform, we can make informed guesses about what kind of evidence our testifiers have without sharing all of it, and we can take account of whether unshared details about their evidence are likely to be significant. Finally, as testifiers, we enjoy the benefit of not having to store and share all of our evidence with our audience of surrogates, which would be enormously costly; a testifier can just report the beliefs she has retained, and forget the rest of the evidence she once based those beliefs on. (It is impossible to store all our evidence anyway, as noted by Harman (1986, ch. 4).)

Reasoning: A group of conformists can also more efficiently accomplish the task of drawing inferences on the basis of known premises. If Wiles reports that Fermat's Theorem is true, or Einstein proposes a creative new theory that best explains the known evidence, we can trust these reports without ourselves discovering or even working through the proposed arguments. As long as we trust that these conclusions are being reached using only applications of belief-forming rules that we accept (applied to known premises), then we don't need to ourselves do the work of applying those rules; we can directly benefit from the work, and the creativity, of others.

Uniqueness Is Required

We have argued that our proposed view can explain the high practical value of promoting rationality. Our view assumed uniqueness. Given uniqueness, promoting rationality involves promoting conformity. Conformity allows us to function as epistemic surrogates, which in turn efficiently ensures the reliability of testimony and enables the division of epistemic labor. Our view thus explains why promoting rationality is an effective and efficient means of meeting our epistemic goal, getting to the truth.

So our explanation works, given uniqueness. And in fact, it *requires* uniqueness. If permissivism is true, then rational reasoners need not conform. That is, there are cases where rational reasoners use alternative belief-forming rules, rules that yield distinct views given the same evidence. In this case, the enforcement of rational rules of reasoning does not make it safe to trust the testimony of rational reasoners, since there is now a risk that a rational reasoner will not be reliable.

Views that deny uniqueness come in many different flavors. Radical permissivist views, which only require consistency or probabilistic coherence (for example), have no hope of explaining our central datum: these views count grue-projectors, counterinductivists, and the like as rational. And promoting mere consistency won't serve our goals, as we discussed above. More moderate views that fall under permissivism may only allow slight differences in rational beliefs given the evidence, or perhaps only a slight difference in rational degree of belief. But, we claim, our view requires rejecting even these moderate views. While these views may not make rational reasoners downright *un*reliable, they still have the result that rational reasoners cannot be ensured to be *as* reliable as they can be if uniqueness is true. Promoting rationality is *less* instrumentally valuable given moderate permissivism than given uniqueness. Thus, uniqueness *best* explains the high value of rationality.

Our view explains why it is worth our while to promote rationality. Having a community of reliable testifiers is beneficial for everyone. And our view captures plausible judgments about what is, in fact, rational: it rules out unreliable epistemic practices like those of a gruesome reasoner or

a committed conspiracy theorist, and it rules out flukey reliable epistemic practices, like clairvoyance. Our view also requires, that is, it entails, the truth of uniqueness. Inferring to the best explanation, then, we infer uniqueness.

Our view is closely related to some views defended in recent literature. In ethics, Gibbard (1990) and Fricker (2016) argue that evaluations of moral praise and blame serve a coordinative function. In epistemology, Dogramaci (2012 & 2015) argues, much as we did, that epistemic evaluations serve a coordinative function. See these authors for further detail and defense of views of the sort we favor. Gibbard and Dogramaci examine and defend many of the empirical psychological and sociological commitments these views carry. Our distinctive interest here, however, is on the implications of such approaches for uniqueness. In a very similar spirit to the present paper, Hedden & Greco (forthcoming) independently explore how the views of Craig, Gibbard, and Dogramaci can be applied to argue in favor of uniqueness.

Let's address some potential worries about our particular view and our use of it to support uniqueness. We'll reply to four objections.

First Objection: An Elitist Permissivism Could Do Just As Well

Objector: You claim that your view requires uniqueness in order to explain your central datum, the value of promoting rationality, and to explain how we can trust one another as epistemic surrogates. But couldn't a certain kind of permissivist say all the same things? I am imagining an elitist sort of permissivist, who forms a community of like-minded friends who share her epistemic rules. She trusts members of her in-group and treats them as epistemic surrogates. But she extends her judgments of *rationality* to some members of the out-group as well as the in-group.

Us: We have three replies to this objection.

First, while it's true that there *could* be a community of permissivists who have a practice like this, we are aiming to explain our *actual* practice, which we do not think works this way. When we are reluctant to accept people's testimony, we are often also reluctant to call them (fully) rational. (You may trust your sister about most things, but not about politics—you say she's been *brainwashed* by Fox News. Your sister doesn't trust you about whether it's safe to drive in this weather—she says you're *paranoid* and *over-cautious*.) Insofar as our practice of epistemic evaluation tends to go hand-in-hand with our practice of trusting one another's testimony, our view provides a better explanation of these actual practices than elitist permissivism does.

Second, this elitist permissivist practice is less efficient than the one we get with uniqueness. If uniqueness is true, we only need to make one kind of evaluation where elitist permissivism would require two. With uniqueness, the *rational* rules are also *trustworthy* rules (i.e. the rules that make their followers

acceptable epistemic surrogates). With this elitist permissivism, we must keep track of the trustworthy rules (which are followed by in-groupers) and the rational rules (which also include those followed by some out-groupers) separately.

Finally, a practice like this makes it mysterious why the in-groupers should bother to give any positive appraisal to the rational out-groupers. Why should we care about promoting (or even identifying) rationality, rather than just in-group rationality? Given uniqueness, there is no such mystery.

Second Objection: The Difficulty of Knowing Which Credence Is Rational

Objector: If, given the evidence, a unique credence in some particular hypothesis is rational, how on earth could we be expected to know what that exact credence is? How could you know, given this particular batch of evidence, that hypothesis *H* requires a credence of exactly 0.53629, or something like that? Or, if you think credences are, or can be, "mushy", and the uniquely rational credence is something like the interval [0.4837–0.683], how could we possibly know this?

Us: We have two points to make in reply.

First, the (moderate) permissivist is in no better position to address these concerns about the epistemology of facts about rationality ("epistemic epistemology", an issue analogous to the issue of so-called "moral episte-mology"). If permissivism is true, then given the evidence, some range of credences, maybe a range of various mushy credences, is rational. How are we able to know what that permitted range is? If anything, it seems *harder* to know.

Second, even though we don't have here a complete account of how we *know* what's rational given evidence *E*, our impermissivist has a neat story about how we can easily and usefully form our views about what's rational given *E*: just take whatever you think *you* would believe if you were to acquire evidence *E*, and consider that view to be the rational view given *E*. To the extent that you are right about what you would believe, promoting that response is what promotes coordination, i.e. what cultivates epistemic surrogates, and we've already seen the value of that.

There are many cases where it's hard to say what the rational response to a body of evidence is, and plausible-looking candidates turn out to be wrong. (It sometimes turns out that what looked like the rational response to our evidence is supported by, e.g., indifference intuitions which turn out to be inconsistent. (See, e.g., Meacham (2014) for discussion of these cases.)) It is sometimes argued that these kinds of cases pose special problems for uniqueness, because they are cases in which it is hard to answer the question of what it is rational to believe. To the extent that it is hard to know how we *would* respond to our evidence in these situations, our proposal will be

admittedly limited. But in spite of that, uniqueness still offers an advantage in conducting epistemic epistemology. Reflecting on what you would believe, given some evidence, is a job that anyone should be able to do, permissivist or not. Insofar as we *can* know what we would believe, our proposal offers the impermissivist, and only the impermissivist, something to say about how we could come to know a good deal about which beliefs are rational.

Third Objection: It Is Too Difficult to Enforce and to Comply with Impermissive Standards

Objector: You've proposed to explain the value of promoting rationality in terms of the practical value of having epistemic surrogates. You claim the ideal surrogate will give testimony based on beliefs that are *exactly* the beliefs the audience would have had themselves if they had acquired and reasoned through the testifier's evidence. But it is very difficult to enforce, and very difficult to meet, such strict standards, and therefore it is an inefficient use of resources to cultivate ideal surrogates, ones with precisely the beliefs we'd have. It is easier, and therefore a more efficient way of pursuing our goals, to cultivate *approximate* surrogates, who have *approximately* the beliefs that we'd have. And this supports moderate permissivism, rather than uniqueness. For example, in a partial-belief framework, a moderate permissivist view might say that given a body of evidence, any credence within some *range* is often permissible—not just one specific credence. In promoting rationality, on this view, we merely want to get one another into the rational range. Since this practice would be more efficient, it provides a better explanation of our practice of rationality ascriptions, and of why that practice is instrumentally valuable.

Us: The objector argues that enforcing and meeting a moderately permissive epistemic standard would be much easier, and from this they infer that a practice of making rationality ascriptions according to a moderately permissive standard would be a more efficient means of pursuing the truth, and thus would be a more instrumentally valuable evaluative practice than an impermissive practice. We disagree. First, we will push back against the claim that a moderately permissive practice would be easier to enforce or easier to meet. Second, we will argue that, even if we had to grant that permissive epistemic standards are easier either to enforce or to meet, it would not follow that a moderately permissive epistemically evaluative practice would help us pursue the truth more efficiently.

First part of our response: would a moderately permissive practice be easier to enforce or meet than an impermissive practice? We grant that, in some cases, relaxing one's standards makes them easier to comply with and easier to enforce. For instance, suppose a company has the choice between two dress codes. The first says that employees must wear khaki pants.

The second says that employees must wear Dockers (R) On-The-Go Marina Straight Fit khakis in Toasted Cardamom. Obviously the second dress code would be harder to implement than the first: it would require careful checking on the part of the enforcers, and careful shopping on the part of the employees.

However, it is not obvious to us that this analogy carries over to the case of epistemic evaluations. First: is a strict standard more difficult to meet, or to enforce, than a permissive one? One way in which a strict standard may be harder to meet is if it is harder to *know* what the standard requires (because it is hard to remember all the specifics, maybe). We can imagine situations where this would be true for dress codes, but have already argued that it is not true for epistemic evaluations. Another way in which a strict standard may be harder to meet is if the means to meet that standard are scarce or hard to come by—for instance, if one has to shop around at multiple stores before finding the right size and model of pants. However, this is not the situation we are in when we "shop around" for doxastic states. In whatever sense beliefs or credences are ever available to us, it seems right to say that they are *all*, *always* available to us. (Horowitz (forthcoming) discusses this point further.) At least, any difficulty we have in conforming our beliefs to the rational requirements (because of doxastic involuntarism, for instance, as well as particular mental pathologies) is not made any *more* difficult if those requirements are impermissive. Similarly, we doubt that impermissive standards make things harder on us as *evaluators*. If it is at least as easy to know what rationality requires given impermissivism as it is given permissivism, it should also be at least as easy to tell whether someone is meeting those standards, given her evidence.

Another motivation for this objection might be the thought that it is especially difficult to enforce *precise* standards, as opposed to vague ones. But even if this is true, it is orthogonal to the issue of permissivism versus impermissivism. The standards allowed by a permissivist are less strict because they allow a *range* of answers, but there is no reason why this range of answers should be any less *precise* than the unique answer enforced by the impermissivist's standard. (Christensen (2007, fn. 8) makes a similar observation.) Similarly, a company dress code that allowed seven different, precise models of khaki pants would not be appreciably easier to enforce, or to comply with, than a dress code that allowed only one. Perhaps plausible epistemic standards should not have precise boundaries and instead allow, somehow, for vague borderline cases. But the permissivist isn't in a better position to accommodate that thought than we are. There are no fewer pesky borderline cases on a view that permits a range of attitudes than there are on a view that permits only one.

Finally, and on a similar note, one might think that it is harder to comply with a rule that requires *sharp* rather than *mushy* credences. But this distinction, too, crosscuts the permissivism/impermissivism distinction.

An impermissivist can happily say that the uniquely rational attitude is a mushy credence, taking on board any risks or rewards that come with mushy attitudes more generally.

So, to sum up the first part of our response to the present objection: we are not convinced that impermissive standards are, because they are impermissive, significantly harder to enforce or to comply with than moderately permissive ones.

Second part of our response: even if we had to grant that a permissive practice *would* be easier to implement, it does not follow that a permissive practice would be overall more instrumentally valuable than an impermissive one. Whether some standard-enforcing practice is instrumentally valuable depends not only on the ease of meeting or enforcing those standards, but also on how well the practice achieves our goals. Indeed, it is often better to have very high standards, and encourage people to approach those high standards—even if those standards are very hard to meet or enforce perfectly, and even if in fact almost nobody will actually meet them. We think that rationality is like this: it makes sense to keep the bar high.

In this respect rationality is not like a company dress code, where more demanding standards are chosen arbitrarily; Dockers (R) On-The-Go Marina Straight Fit in Toasted Cardamom are not *better* than (most) other khakis. There *is* something better about *our own* epistemic standards; they recommend the most accurate, or most reliable, beliefs (as compared to any competing views a permissivist might want to also allow). This means that our own standards will get us the most reliable testimony from others. In bringing about conformity to our own standards, we are improving the accuracy and reliability of our surrogates.

(See Horowitz (2013) and Schoenfield (2014) for further discussion of this issue in the context of permissivism. According to Horowitz (2013), because we see our own epistemic standards as the most reliable standards available, we should prefer that others have our standards as well. Schoenfield (2014), however, argues that we may coherently regard our own standards are more *accurate* than others, but accept other standards as equally *rational*. We disagree with Schoenfield: endorsing some standards as rational involves encouraging others to take up those standards. And to ensure the reliability of testimony, we should encourage others to take up the most reliable standards available: our own.)

As a better analogy than dress codes, consider the rubric a teacher uses to grade papers. When you are grading, you hold students to a high standard. Almost nobody gets an A+. Perhaps the whole grading process would be *easier* if you lowered your standards: it would be easier to *get* an A+, and perhaps easier to assign grades as well. But even granting this, an easier grading scheme would be much *less* instrumentally valuable, overall, than a strict grading scheme. That's because an easier grading scheme is much less effective at achieving a very important goal, to end up with a good batch

of papers at the end of the semester. Indeed, part of the tragedy of grade inflation is that it takes away from this very important instrumental use of grades.

Holding one another to high epistemic standards is, we have argued, quite instrumentally valuable. When others use the same epistemic rules that we do, they can be our epistemic surrogates. But to the extent that their rules differ from ours, their testimony becomes less trustworthy. Perfect conformity to one set of epistemic rules is a *demanding* standard: it is hard to meet, but it is worth shooting for.

Fourth Objection: Trusting Epistemic Surrogates with Different Jamesian Values

Objector: Your argument assumes that it only makes sense to trust epistemic surrogates who believe as you would, given their evidence, because these surrogates are the ones you can most efficiently take to be reliable sources of true belief. But this assumption papers over an important point. There are different ways to assess the truth and reliability of someone's beliefs and epistemic methods, corresponding to different trade-offs between the two Jamesian goals: *believe truth* and *avoid error*. Suppose there are many rationally permissible ways to weigh those two goals against one another. Then suppose that those different weightings also lead to different rational responses to a single body of evidence. (Horowitz (forthcoming) argues that this supposition is false; but let's grant it for now.) It would then make sense for you to evaluate others as rational, just as long as their weighting of the Jamesian goals was permissible—even if those others don't believe as you would, given their evidence.

Kelly (2013, section 2) suggests this Jamesian picture as a possible motivation for permissivism. He writes: "I think that someone who arrives at the conclusion that Uniqueness is false in this [Jamesian] way should not feel especially threatened by the kinds of arguments offered by Roger [White]". That is: a Jamesian permissivist who denies the in*ter*personal thesis, uniqueness, should not feel especially threatened by White's (2005 & 2013) arguments, since those only support the in*tra*personal thesis, uniqueness*.

Us: We disagree with Kelly: the permissivist should feel threatened. When supplemented with our picture of the value of promoting rationality, White's arguments cause the Jamesian objection to uniqueness to fail.

The objector is right that on our view, regarding someone else's beliefs as *rational* involves regarding her as a reliable epistemic surrogate. But this does not mean merely assessing their beliefs as "reliable" or "truth-conducive" according to *some* sensible way of measuring reliability or truth-conduciveness; it means being ready to trust their testimony, and to take on their beliefs without having to sort through all of their evidence yourself. If this is right, you

should not regard the beliefs of someone else who strikes a different trade-off between the Jamesian goals as rational beliefs unless you, yourself, are willing to make that trade-off. (If you were unwilling to make that trade-off yourself, you would also be unwilling to take the person's testimony at face value.) Now suppose that the second step of the Jamesian argument is right, and that different Jamesian trade-offs license different beliefs in response to a single body of evidence. Then, assessing more than one set of Jamesian values as rational means being willing to take on for yourself the beliefs that each of those values license. *Interpersonal* permissivism requires *intrapersonal* permissivism as well.

It is precisely this feature of intrapersonal permissivism—being willing to adopt several different responses to a single body of evidence—that White's arguments target. (White argues for (among other things) the irrationality of "flip-flopping" between two sets of rational beliefs, and of treating rational beliefs other than your own as suitable for practical deliberation.) Our view draws a connection between the intrapersonal and interpersonal theses; it provides reason to adopt uniqueness, given uniqueness*.

Interestingly, many defenders of permissivism are happy to deny the intrapersonal version, and to grant uniqueness*. And permissivist critiques of White's arguments have been largely focused on the fact that they seem to support uniqueness* rather than uniqueness. (See, for instance, Kelly (2013), Schoenfield (2014) and Meacham (2014) on both of these points.) Many permissivists should, therefore, be happy to take White's arguments on board. But if we are right, accepting White's arguments also means rejecting the Jamesian objection to uniqueness. If you are not willing to accept testimony from those with different Jamesian values, you should also be unwilling to evaluate them as epistemically rational.

Acknowledgements

References

Ballantyne, N. & Coffman, E.J. (2011), "Uniqueness, Evidence, and Rationality", *Philosophers' Imprint* 11 (18): 1–13.

Berker, S. (2013), "Epistemic Teleology and the Separateness of Propositions", *The Philosophical Review*.

Boghossian, P. (2007), "Epistemic Rules", *The Journal of Philosophy* 105 (9): 472–500.

BonJour, L. (1980), "Externalist Theories of Empirical Knowledge", *Midwest Studies in Philosophy* 5 (1): 53–70.

Christensen, D. (2007), "Epistemology of Disagreement", *The Philosophical Review* 16 (2): 187–217.

Craig, E. (1990), *Knowledge and the State of Nature*, Oxford University Press.

Dogramaci, S. (2012), "Reverse Engineering Epistemic Evaluations", *Philosophy and Phenomenological Research* 84 (3): 513–530.

———. (2015), "Communist Conventions for Deductive Reasoning", *Noûs* 49 (4): 776–799.

Feldman, R. (2007), "Reasonable Religious Disagreement", in *Philosophers without God*, ed. Antony, L., Oxford University Press.

Fricker, M. (2016), "What's the Point of Blame?", *Noûs* 50 (1): 165–183.

Gibbard, A. (1990), *Wise Choices, Apt Feelings*, Harvard University Press.

Goldman, A. (1986), *Epistemology and Cognition*, MIT Press.

———. (1992), "Epistemic Folkways and Scientific Epistemology", in his *Liasons*, MIT Press.

Greco, D. & Hedden, B. (forthcoming), "Uniqueness and Metaepistemology", *Journal of Philosophy*.

Harman, G. (1986), *Change in View*, MIT Press.

Horwich, P. (1990/98), *Truth*, Oxford University Press.

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

———. (forthcoming), "Epistemic Value and the 'Jamesian Goals'", in *Epistemic Consequentialism*, eds. Dunn & Ahlstrom-Vij, K., Oxford University Press.

Kelly, T. (2003), "Epistemic Rationality as Instrumental Rationality: A Critique", *Philosophy and Phenomenological Research* 66 (3): 612–640.

———. (2010), "Peer Disagreement and Higher Order Evidence, in *Disagreement*, eds. Feldman, R. & Warfield, T., Oxford University Press.

———. (2013), "How to Be an Epistemic Permissivist", in *Contemporary Debates in Epistemology*, eds. Steup, M. & Sosa, E., Wiley-Blackwell.

Lehrer, K. (1990), *Theory of Knowledge*, Westview.

Lewis, D. (1974), "Radical Interpretation", *Synthese* 27(3/4): 331–344.

Madison, B. (forthcoming), "Epistemic Value and the New Evil Demon", *Pacific Philosophical Quarterly*.

Meacham, C. (2014), "Impermissive Bayesianism", *Erkenntnis* 79 (6): 1185–1217.

Plantinga, A. (1993), *Warrant: The Current Debate*, Oxford University Press.

Pollock, J. & Cruz, J. (1999), *Contemporary Theories of Knowledge*, Rowman and Littlefield.

Rosen, G. (2001), "Nominalism, Naturalism, Epistemic Relativism", *Philosophical Perspectives* 15: 69–91.

Schoenfield, M. (2014), "Permission to Believe", *Noûs* 48 (2): 193–218.

Stalnaker, R. (1984), *Inquiry*, MIT Press.

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

———. (2013), "Evidence is not Permissive", in *Contemporary Debates in Epistemology*, eds. Steup, M. & Sosa, E., Wiley-Blackwell.

Williamson, T. (2000), *Knowledge and its Limits*, Oxford University Press.

———. (2009), "Reply to Hawthorne and Lasonen-Aarnio", in *Williamson on Knowledge*, eds. Greenough & Pritchard D., Oxford University Press.