

Do Great Minds Really Think Alike?

Christopher Willard-Kyle

January 8, 2016

Abstract

Recently, a number of epistemologists (notably Feldman [2007], [2009], and White [2005], [2013]) have argued for the rational uniqueness thesis, the principle that any set of evidence permits only one rationally acceptable attitude toward a given proposition. In contrast, this paper argues for extreme rational permissivism, the view that two agents with the same evidence (evidential peers) may sometimes arrive at contradictory beliefs rationally. This paper (1) identifies different versions of uniqueness and permissivism that vary in strength and range, (2) argues that evidential peers with different interests need not rationally endorse all the same hypotheses, (3) argues that evidential peers who weigh the theoretic virtues differently (that is, who have different standards) can sometimes rationally endorse contradictory conclusions, and finally (4) defends the permissivist appeal to standards against objections in the works of Feldman and White.

(N.B. This is the accepted manuscript of “Do Great Minds Really Think Alike?” published in *Synthese*, January 2014. The published version can be found at <http://link.springer.com/article/10.1007/s11229-015-0984-x>.)

1 Introducing Uniqueness and Permissivism

1.1 Rational Uniqueness

It is proverbial wisdom that great minds think alike. And like most proverbs, this saying communicates an important insight. Often there is only one correct solution to a problem and many ways to get it wrong. And especially when the correct solution requires a special measure of insight or intelligence to see, it is unsurprising that the great minds are the ones who arrive at a common solution. Some epistemologists think that, given appropriate qualifications, it is more than mere proverbial wisdom that great minds think alike: it is also an important epistemic principle. Such epistemologists believe that whenever two perfectly rational agents—two great minds—have the same evidence, they must of necessity come to the same conclusions.

Is this so? Do rational people who have the same evidence always, in the end, come to the same conclusions? Or is rational disagreement possible even when no mistakes of reasoning have been made? This is the question of rational uniqueness.

(RUE) Rational Uniqueness plus Evidentialism For all complete sets of evidence E , for all epistemic agents S , and for all propositions p , if S has

E, then there is just one doxastic attitude A such that it is rationally permissible for S to bear A toward p.

In other words, rational uniqueness is the thesis that the rationality of doxastic attitudes is a function of our evidence. It is not a mere relation of our evidence, nor is it a function of our evidence plus something else. The function of rationality takes only evidence as its input and results in only one doxastic attitude as its output.

What are the domain and range of this function? The domain includes all possible sets of total evidence. It is important that the input is *total* evidence, since the impact of any *piece* of one's evidence (however such evidence is individuated) depends on the rest of one's evidence. Presumably there could be other kinds of rational uniqueness that are not evidence-based. Such theses might take the following form:

(RU) Rational Uniqueness: For all rational bases B, and for all epistemic agents S, and for all propositions p, if S has B, then there is just one doxastic attitude A such that it is rationally permissible for S to bear A toward p.

In this formula, “rational bases” stands in for whatever it is that fixes what is rationally permissible. Unlike (RUE), it does not specify that evidence exhausts the rational bases. Despite the plurality of potential uniqueness principles, it is predominantly the evidentialist version that has been defended in the literature (See White [2005], [2013], and Feldman [2007], [2009].) This paper shall similarly focus on evidentialist versions of uniqueness. It is worth noting, however, that rejecting evidential versions of uniqueness might turn out to be compatible with endorsing non-evidential (or, at any rate, *extra-evidential*) versions of uniqueness.

The appropriate range of (RUE) is trickier. Are doxastic attitudes best parsed in terms of beliefs? If so, there are only three candidates for doxastic attitudes: believing, disbelieving, and abstaining from belief. Or are doxastic attitudes best parsed in terms of credences? In this case, there are infinite doxastic attitudes that correspond to each number between 0 and 1. If doxastic attitudes are best expressed in terms of credences, are the outputs always point-sized, or do sets of evidence sometimes output an extended interval as the rational attitude?¹² This paper does not take a stand on such issues. The arguments I will present aim to be successful regardless of what type of attitude the rational uniuquist accepts as fundamental.

A related and difficult question is whether the range should include, in addition to doxastic attitudes, absences of doxastic attitudes. Suppose someone has evidence that supports belief that p. Rational uniqueness entails that she should not, upon considering her evidence, decide to withhold belief about p. Does uniqueness also entail that it would be impermissible for her never to form any judgment concerning p in the first place? As articulated in the definition

¹On this picture, it would be rationally required to have, say, a credence expressed as the interval [.5, 6.1), and it would be impermissible to have a credence that is more specific than that interval.

²The various kinds of doxastic attitudes and their applicability to the uniqueness and disagreement literatures are nicely summarized in Kvanvig's (2014) *Rationality and Reflection* (pp. 85).

of (RUE) above, the answer is “no.” After all, never having formed an opinion about p is different than considering the evidence for p and then deciding to withhold judgment. Never having formed an opinion about p is not some fourth kind of doxastic attitude one could bear toward p , but the absence of any doxastic attitude regarding p whatsoever.

Matters are complicated, however. Consider an analogue in ethics. It is natural to think of the output of the relation that expresses moral permissibility as ranging over possible actions. We make judgments such as, “It is permissible help a friend in a crisis,” or “It is impermissible to make sport of one’s friend when she is in trouble.” Nevertheless, we sometimes ascribe the same language of moral permissibility or impermissibility to non-actions too, as in the sentence, “It is impermissible to do nothing when an innocent person is drowning nearby when one can help at low cost.” Someone might object that what we really object to in this case is the *decision* to do nothing. And deciding—even deciding to do nothing—is a kind of activity. Still, I think that in many cases we would blame someone for doing nothing even when no decision has been made. Suppose the passerby was so callous that the thought of helping the drowning victim didn’t even occur to her (and therefore that she made no decision about whether or not to help). I think we would still say that her non-action was impermissible. She is guilty of negligence.

There is reason to think, therefore, that in ethics the range of things to which we ascribe moral permissibility includes not just actions but also the absence of actions, even when those absences are unintentional. We condemn some instances of inactivity as negligent, but permit others. This raises the question of whether the evaluation of rational permission, analogously to the evaluation of moral permission, should range over absences of doxastic attitudes as well as doxastic attitudes themselves. Moreover, it raises the question of whether the best version of the rational uniqueness thesis should range over only doxastic attitudes or both doxastic attitudes and their absences.

Perhaps it is best to think of there being two different uniqueness principles. The first, (RUE), which we already articulated, is repeated below:

(RUE) Rational Uniqueness plus Evidentialism: For all complete sets of evidence E , for all epistemic agents S , and for all propositions p , if S has E , then there is just one doxastic attitude A such that it is rationally permissible for S to bear A toward p .

A more expansive form of uniqueness ranges over both doxastic attitudes and their absences. Let us introduce the term “doxastic state” to signify either a doxastic attitude or the absence of a doxastic attitude.

(E-RUE) Expansive Rational Uniqueness plus Evidentialism: For all complete sets of evidence E , for all epistemic agents S , and for all propositions p , if S has E , then there is just one doxastic state D such that it is rationally permissible for S to bear D toward p .

According to both (RUE) and (E-RUE), rational permissibility is a function of one’s evidence. But (E-RUE), unlike (RUE), includes absences of doxastic attitudes within the range of the function. One way of articulating the difference between these two principles is to say that (E-RUE) takes the possibility of rational negligence more seriously; for whenever a positive doxastic attitude

toward a proposition is permissible, (E-RUE) entails that failing to adopt any attitude toward it is forbidden. On (E-RUE) but not (RUE), the rationally permissible and rationally obligatory overlap completely. Though (E-RUE) is the more demanding principle, it does not logically entail (RUE). In fact, if (E-RUE) is true and it is ever the case in which it is rationally permissible to have no doxastic attitude toward p , then (RUE) must be false, for then there would be a case in which zero (and so not one) doxastic attitudes would be rationally permissible.

We can cook up even stronger principles than (E-RUE). For instance, an incredibly strong version of uniqueness requires that one doxastic attitude toward every proposition is not only permitted but obligatory given some evidence:

(S-RUE) Super-Strong Rational Uniqueness: For all complete sets of evidence E , for all epistemic agents S , and for all propositions p , if S has E , then there is one doxastic attitude A such that it is rationally obligatory for S to bear A toward p .

I doubt whether anyone has ever held (S-RUE). It is very strong, requiring that one should have a single, positive doxastic attitude toward every proposition whatsoever.³ Still, it is worth noting the various strengths, domains, and ranges that uniqueness principles can encompass. (E-RUE) and especially (RUE) look relatively modest in comparison. No distinct arguments in this paper will be given against (S-RUE) specifically; however, on the plausible assumption that rational obligation entails rational permission, (S-RUE) logically entails both (RUE) and (E-RUE). Accordingly, all arguments against either (E-RUE) or (RUE) can double as arguments against (S-RUE).

Most defenders of uniqueness are best interpreted as defending (RUE) rather than (E-RUE). For instance, Richard Feldman, one of the foremost defenders of uniqueness, explicitly denies a logical consequence of (E-RUE) in his excellent article “The Ethics of Belief” (2000: 678-80).⁴ In keeping with the literature, the majority of this paper shall focus on arguments for and against (RUE). Both (RUE) and (E-RUE), however, are worthy of consideration, as engaging with each provides opportunity to explore different ways that maximally rational peers may or may not permissibly arrive at different conclusions.⁵

It is worth noting that evidentialism does not by itself entail any version of uniqueness. Rather, evidentialism entails the following principle of symmetry:

³Feldman considers a principle like (S-RUE) that he labels (O1):

“For any proposition p , time t , and person S , S epistemically ought to have at t the attitude toward p that is supported by S ’s evidence at t ” (Feldman, 2000: 678).

Although he toys with the idea that (O1) might be correct if such attitudes are interpreted as merely dispositional, he ultimately discards (O1) in favor of a principle like (RUE), namely (O2):

“For any person S , time t , and proposition p , if S has any doxastic attitude at all toward p at t and S ’s evidence at t supports p , then S epistemically ought to have the attitude toward p supported by S ’s evidence at t ” (Feldman, 2000: 679).

⁴Feldman claims that “no attitude is epistemically required, but only one is epistemically permitted” (Feldman, 2000: 680). In other words, Feldman thinks it is both permissible not to have any attitude toward p and permissible to have the (lone) permitted doxastic attitude toward p . This means that two different doxastic states are permitted, thus contradicting (E-RUE).

⁵I am thankful to Andrew Moon and an anonymous reviewer for helping me to appreciate the important differences between (RUE), (E-RUE), and (S-RUE).

(Sym) If epistemic agents S and T have the same complete evidence, then it is rationally permissible for S to bear doxastic attitude A toward p iff it is rationally permissible for T to bear A toward p.⁶

(Sym) says nothing about *how many* doxastic attitudes are rationally available to an agent given the evidence, only that the same and only the same options are available to anyone else who has the same evidence.

In opposition to uniqueness, some philosophers endorse rational permissivism (or “optionalism” as in Kvanvig [2014]). Permissivists believe that, at least sometimes,⁷ multiple doxastic attitudes are permissible relative to the same evidence.⁸ It would be tempting to define permissivism as the negation of (RUE) and (E-RUE); however, as Matthew Lee has reminded us, permissivism is only a contrary of uniqueness and not a contradictory (Lee, 2013). Suppose that genuine rational paradoxes are possible in the following way: in some epistemic situations, no matter what one does or fails to do, one will violate the rational norms. In addition, suppose that no more than one doxastic attitude is ever permitted (and so that permissivism is false). If this is the way things are, then (RUE) is false, for it is possibly the case that not even one doxastic attitude is rationally permitted given one’s evidence. Similarly (E-RUE) is false, for it is possibly the case that no attitudes *or* absences of attitudes are rationally permitted. It is conceivable, therefore, that neither rational permissivism nor any form of uniqueness is true.⁹

So (RUE), (E-RUE), and rational permissivism do not exhaust the conceptual landscape. Nevertheless, it is doubtful whether the sort of rational paradoxes required to avoid both uniqueness and permissivism are really possible. This paper will assume that there is always something we can do (or refrain from doing) without violating the rational norms. At any rate, all of the arguments I give against uniqueness could also be given as arguments against the thesis that *at most one* doxastic attitude is rationally permissible for an agent given her evidence. Thus we may, for all practical purposes, treat rational permissivism as the negation of (RUE) and (E-RUE). Rational permissivism is true just in case that both (RUE) and (E-RUE) are false.

In what follows, I argue for rational permissivism and against rational uniqueness. It is possible for two people with the same evidence—let us call such “evidential peers”¹⁰—to disagree about a proposition without either one having made a mistake in reasoning. The question of uniqueness (How many doxastic attitudes toward a given proposition can be permitted by a total body of evidence?) has often been explored in connection to questions of disagreement

⁶As with uniqueness, we could also consider an expanded version of symmetry that ranges over all doxastic states (including absences of attitudes).

⁷Kvanvig (2014) helpfully distinguishes between synchronic and diachronic versions of uniqueness and permissivism (or “restrictivism” and “optionalism”). All the versions of uniqueness and permissivism in this paper should be understood as synchronic theses.

⁸This does not necessarily mean that multiple doxastic attitudes are permissible for the same person.

⁹Rational permissivism, unlike either formulation of uniqueness, is compatible with genuine rational paradoxes, for permissivism only claims that multiple doxastic attitudes are sometimes permitted, not that they always are.

¹⁰Evidential peers, like epistemic peers, have all the same evidence. They need not be equally competent. Competence is not strictly relevant to the question of whether two people can rationally come to different conclusions on the same evidence. Even fools believe rationally sometimes.

(How ought we to respond when we find that epistemic equals disagree with us?). And while the former is certainly important for the latter, it is surely also of intrinsic interest whether evidence is ever permissive. Here I will confine myself to the former issue. I am not, for instance, defending the stronger claim that epistemic peers can continue to disagree after their disagreement becomes manifest, only that evidential peers can (perhaps unknowingly) disagree rationally.

Sometimes great minds think very differently.

1.2 Rational Permissivisms

It is a notable feature of epistemology that we often apply normative language to doxastic attitudes. We speak of inferences as being either *good* or *bad*. We say things like, “You had no *right* to assume that!” when blaming someone for having a belief, or “You *should have* known!” when blaming someone for not having a belief. As epistemologists, we care greatly whether we are *justified* in our beliefs. When we believe something true, we are *right*, and when we believe something false, we are *wrong*. Let the sentential operator “ \Box_R ” mean “rationally obligatory.” Taking deontic and modal logic as guides, the following schema is plausible where Bp is belief that p is true:¹¹

Rationally Obligatory: $\Box_R Bp$

Rationally Impermissible: $\Box_R \sim Bp$

Rationally Optional: $\sim \Box_R Bp \wedge \sim \Box_R \sim Bp$

These categories are exhaustive and exclusive. All beliefs fall into one and only one of these three categories. The rationally optional and rationally obligatory together are the rationally permissible beliefs:

Rationally Permissible: $\sim \Box_R \sim Bp$

(E-RUE) entails that, relative to some evidence, there are no rationally optional beliefs; there is no set of evidence such that for any proposition p , $\sim \Box_R Bp \wedge \sim \Box_R \sim Bp$. (RUE) does not entail that there are no optional beliefs, as it may be permissible never to have formed a judgment about a proposition. It is rationally obligatory, however, that conditional on forming a judgment about p , one endorses the uniquely rational attitude. Permissivists disagree. Some permissivists are content to claim that some beliefs are rationally optional even given that one has formed a judgment about a proposition. Many permissivists, however, also want to say that it is sometimes permissible for one person to believe that p while another person with the same evidence permissibly believes that $\sim p$. Even given that an agent has formed a judgment about p , there are some sets of evidence such that there are some p such that $(\sim \Box_R Bp \wedge \sim \Box_R \sim Bp) \wedge (\sim \Box_R B \sim p \wedge \sim \Box_R \sim B \sim p)$.

Let us distinguish between what we may term “weak” and “strong” disagreement. Weak disagreement occurs when one person believes a proposition and another person withholds belief from that same proposition. Strong disagreement occurs when one person believes a proposition and another person believes the negation of that proposition. The theist and the atheist strongly disagree, whereas the agnostic and the theist (or the atheist) disagree weakly. Whether a

¹¹This model does not require doxastic voluntarism. We can describe beliefs (or other doxastic attitudes) as obligatory, permissible, etc., even if they are actions we cannot voluntarily perform.

disagreement is weak or strong in this sense has nothing to do with the intensity of the disagreement.

Among weak disagreements, we must be aware of another distinction. There are disagreements of commitment and disagreements of incidence. A committed disagreement about p is one in which both parties have considered whether or not p . In an incidental disagreement, one party has not reflected about whether or not p . (If neither party has considered whether or not p , there can be no disagreement). Consider again the agnostic. There are really (at least) two types of agnostics. First, there is the committed agnostic who, reflecting on the evidence, neither comes to believe that God exists nor that God does not. In contrast, the incidental agnostic simply hasn't fully considered the question of God's existence.¹² Perhaps she has grown up in a context where the subject is rarely if ever brought up. Whether a disagreement is committed or incidental in this sense has nothing to do with how earnestly one is living out her convictions. Accordingly, there cannot be an incidental theist or an incidental atheist in this sense. If someone believes p or not- p , she must have considered p , even if briefly.¹³

These definitions have been presented in terms of beliefs. If someone prefers to think of doxastic attitudes as credences, however, we may make similar distinctions. Weak and strong disagreements would be defined by the relative disparity between credences. If one person assigns a credence of .98 and another of .97, they have a relatively weak disagreement. If the second credence were .41, however, it would constitute a relatively strong disagreement. Just how disparate credences must be to be considered a "strong" disagreement may be contextually determined. Or if there is no firm boundary between "strong" and "weak" disagreement, we can still speak of disagreements as "stronger" or "weaker." An incidental disagreement occurs when one party fails to assign a credence concerning p .

As defined, rational permissivism says nothing about what kind of rational disagreement is possible. But most interesting cases of disagreement involve not only weak but also strong disagreement, not only incidental disagreements, but also disagreements of commitment. Philosophical, political, and moral disagreements are often both strong and committed. Some believe that physician-assisted suicide is permissible, but others who seem to have access to the same evidence believe it is forbidden; some believe one candidate should be elected, but others believe her opponent should be elected instead; some believe that there are abstract universals, whereas others believe there are no such things.

Earlier we claimed that rational permissivism is true if both (RUE) and (E-RUE) are false. Let us now distinguish between moderate and extreme rational permissivism. Moderate permissivists simply believe that (RUE) and (E-RUE) are false.¹⁴ I will argue for what I shall call "extreme rational permissivism." Extreme rational permissivism holds not only that (RUE) and (E-RUE) are false, but also that there are, or at least can be, cases of rational, strong, and committed disagreements among evidential peers.

We now have the conceptual tools and framework we need to approach the

¹²Mark Nelson helped me fully see this distinction.

¹³At least it seems so at a first glance, although perhaps certain kinds of implicit beliefs can be maintained without ever considering them.

¹⁴Or, better, moderate permissivists simply believe that (RUE) and (E-RUE) are false and are not false solely because rational paradoxes are possible.

debate over rational uniqueness. In what follows, I argue against uniqueness and for extreme rational permissivism. In Section 2 I will argue that (E-RUE) is false: it's sometimes permissible not to have any attitude toward a proposition even though at least one positive doxastic attitude toward it is clearly permitted. In Section 3 I will construct a positive case for extreme rational permissivism and defend it against objections that extreme permissivism is *too* extreme. In Section 4 I will consider the best arguments in the literature for uniqueness (White [2005], [2013], Feldman [2007], [2009]) and develop new permissivist responses. What we find will not only suggest that uniqueness is false, but also give us reason to suppose that there are strong, rational, committed disagreements among evidential peers.

2 The Argument from Efficient Evidence: The Falsity of (E-RUE)

In his intriguing article “We Have No Positive Epistemic Duties,” Mark Nelson argues just that. Neither Nelson nor anyone else has applied his arguments to the debate surrounding rational disagreement; nevertheless, if he is right, (E-RUE) is surely false. According to (E-RUE), the epistemically permissible and obligatory overlap completely, and so if we have no positive epistemic duties, neither do we have any positive epistemic permissions. And surely this would be a most repugnant skepticism, such that no positive beliefs (that I have a hand, that $2+2=4$, etc.) are even permissible.

I will not argue the strong case that we have no positive epistemic duties. But through an exploration of Nelson’s argument, I will argue that, in some situations, we have no positive epistemic duties to believe any particular proposition and that some of these are cases in which some positive doxastic attitudes are clearly permissible. As we shall see, the cases Nelson presents as paradigmatic seem to be clear cases of rational, weak disagreement, and thus counterexamples to (E-RUE).

According to Nelson, “We have negative epistemic duties, but no positive ones. There are things that we ought not to believe, but there is nothing that we ought to believe, on purely epistemic grounds” (Nelson, 2010: 83). The qualification “on purely epistemic grounds” is vital, for there are many times when we ought to believe something for non-epistemic reasons. If I want to finish a set of math problems, I ought to know the answers; if I am a Catholic priest, I ought to believe that God exists; if I am sufficiently interested in philosophy, I ought to know about Plato’s theory of forms. But none of these obligations is purely epistemic. They arise because of other pragmatic, ethical, or disciplinary norms.

Nelson’s argument relies on what he calls the “infinite justificational fecundity of evidence,” the idea that “every single bit of evidence ...potentially epistemically justifies an infinitely large array of different beliefs” (Nelson, 2010: 96). In a trivial sort of way, this is quite clearly the case for any evidence that justifies belief in at least one proposition p . For any evidence that justifies p also justifies $q \supset p$, $r \supset p$, $\sim\sim p$, $\sim\sim\sim p$, etc... But Nelson does not merely appeal to logical manipulation of the evidence. Another aspect of the justificational fecundity of evidence is the variety in levels of generality and specificity

licensed by the evidence. Nelson gives the following example:

Given the appearance of some distinctive dark, winged shapes, moving across my visual field, what should I believe? That visual evidence joined with other factors, may license me to believe propositions such as: (1) There are things moving through the air in front of me, (2) There are birds flying in front of me, (3) There are jackdaws flying in front of me, (4) At least three jackdaws exists. (Nelson, 2010: 87)

On this view, there are negative epistemic obligations not to believe anything that contradicts the licensed beliefs (Nelson, 2010: 87). But, there isn't an obligation to believe any particular belief among the set of licensed beliefs. If there is a sense in which one is obligated to endorse any of the licensed beliefs, it is only given non-epistemic factors such as one's interests. If I am interested in the question of how many jackdaws there are in the world, perhaps I should—in some broader sense of “should”—form a belief that (4), whereas “if ...I am merely looking to hail a taxi, I need not form any of those beliefs. I do nothing wrong, epistemically speaking, if I *do* form such beliefs, but, equally, I do nothing wrong if I do not” (Nelson, 2010: 88).

Nelson cites as an advantage of his position that it is psychologically feasible, drily noting just how busy he should have to be in order to believe everything that is licensed for him (Nelson 2010: 86). This is an important advantage; nevertheless, the advocate of (E-RUE) might try to bite the bullet here. Perhaps the norm of rationality is not feasibly attainable, but this by itself hardly means it cannot be a genuine norm. Very few objects in the actual world fulfill the conditions for being absolutely straight, but this does not mean that we should lower our standards for what counts as being absolutely straight. Rather, we recognize that being straight is an ideal property, rarely realized by non-mathematical objects. When we say that a road is straight, we usually mean that it is straight more or less. It is relatively straight, though it may have a few small turns or digressions. Perhaps we mean something similar when we talk about rational behavior. Perhaps people can act in more or less rational ways even if absolutely rational behavior is very difficult. They are “rational” insofar as they approximate true rationality.

I think this is a dangerous line for the advocate of (E-RUE) to take. First, there is an important disanalogy in that finding a straight road is very difficult, but being ideally rational in the sense outlined above—more than merely difficult—is impossible for finite beings. It would require an infinite mind to contemplate every proposition that is supported by the evidence. God might be rational on this interpretation, but the term would be useless as applied to humans. As Scott Stapleford surmises, “if the upshot of a normative theory is that every agent is always and everywhere, necessarily guilty of violating an infinite number of duties—if the rules of the game make success in principle impossible and guarantee failure across the board—then one has good reason to suspect that something has gone wrong with the theory” (Stapleford, 2013: 4068). Second, if rationality is merely approximation of an unattainable ideal, it isn't obvious why we should care whether or not rational disagreement is possible among evidential peers. After all, *nearly* rational disagreement is almost certainly possible. If the only rationality we can attain is near rationality

anyway, why should we care whether or not perfect rationality excludes disagreement? Rational uniqueness is preserved only by dispensing with actual examples of full rationality. It would be best for the defender of (E-RUE) to look for objections elsewhere.

The psychological feasibility of permissivism is a strong point in its favor; however, Nelson is not content to rest his case on the observation that the alternative seems excessively demanding. Imagine an ideally rational observer is looking through a telescope. Suppose we know exactly what the view from the telescope is (i.e. the total perceptual evidence) and all of the observer's background beliefs. Could we thereby come to predict what this observer would believe on the basis of this information alone?

It seems not. Imagine that a telescope displays "distinctive dark, winged shapes, moving across [the observer's] visual field." We cannot predict which of propositions (1)-(4), if any, she will believe. She will probably only form a belief in (4) if she is interested in counting jackdaws. If she is listlessly looking through the lens only to escape boredom, she may not infer any of them. We cannot tell which propositions the ideal observer will infer without knowing, in addition to her evidence: her interests.¹⁵ The fact that we cannot predict what she will believe, even though we know all her evidence and assume her perfect rationality, seems to indicate that rationality does not require the formation of any particular belief (1)-(4).

Someone will say, "Fair enough. Let's say Nelson's examples succeed when considering perceptual evidence. After all, it does seem impossible to predict what an ideally rational observer would believe just by knowing what she sees when she looks around. But why think that that sort of evidence is relevant for (E-RUE)? Perhaps the only evidence that matters for (E-RUE) is propositional evidence."¹⁶ This response ignores the force of the infinite number of logical truths that are justified by any propositional evidence, but let us set these aside for a moment. Even without appealing to logical addition or double negation, we can think of cases similar to Nelson's by looking only at propositional evidence. Presumably, there are many truths of mathematics for which I know all the premises necessary for generating the conclusion, and all the logical rules by which I could arrive at said conclusion, but which I nevertheless fail to know. Am I irrational in this? Certainly it would be irrational to believe that they were false, but have I really violated the epistemic norms simply because I haven't taken the time and energy to discover all of these mathematical truths? This case can be formulated using only propositional evidence and thus remain a problem for (E-RUE).

Someone else will object, "But there are positive epistemic duties! For instance, we ought to believe the truth!" There is a short answer and a long answer to this objection. The short answer is that, unlike Nelson, I need not defend the claim that there are never *any* positive epistemic duties, just that there are *some* situations in which we do not have positive epistemic duties though positive doxastic attitudes are obviously permitted for us.

The long answer is that the sort of duties mentioned earlier are best under-

¹⁵See Nelson 2010, 97-101 for further development of this argument. I owe thanks to Mark Nelson for patiently drawing my attention to the distinction between this argument and the argument from demandingness.

¹⁶I leave aside questions about the nature of evidence and whether all evidence is ultimately propositional.

stood, as Stapleford has convincingly argued, as imperfect duties. Imperfect epistemic duties often require some positive doxastic attitudes although they do not require belief in any one proposition (Stapleford, 2013). Consider imperatives such as, “believe the truth,” or as in Stapleford, “deliberate on [your] evidence” (Stapleford, 2013: 4069). There is a weak (imperfect) and a strong (perfect) reading of this demand. The strong reading requires that epistemic agents believe all truths. The weaker reading suggests that they must know some of them.

An analogy to an imperfect moral duty may be illustrative. Consider the ethical obligation to be generous. In most cases, there are several ways to fulfill this obligation. I could be generous by giving money to a charity, or to my children’s education, or by being generous with my time in a maligned community. In fact, there are so many ways of being generous, that I couldn’t possibly do all of them. A plausible thing to say in this case is that though I have an obligation to do some acts that are generous, I am not obligated to do any particular one of them. Or if I am required to perform one of them, the reason for the obligation will not be solely the duty to be generous. Supposing I succeed at being generous, the explanation for why I do any particular generous act will probably not merely be that I was trying to fulfill my obligation of being generous, for I could have fulfilled this obligation in many other ways. My reason for choosing that particular generous act will probably include my interest in the beneficiary, my contextual circumstances, or other non-moral factors. The same may be true in the epistemic case. It may be obligatory that we believe some truths (know some truths, exhibit some epistemic virtues, etc. . . .) without being obligatory that we believe every, or even any particular, one of them.¹⁷

How can we tell whether our duties to be rational are imperfect duties in the same way that our duty to be generous is imperfect? Stapleford notes that imperfect duties are often marked by the ability to excuse oneself from a failure to satisfy the duty in one instance by one’s success in fulfilling the same duty in another, whereas perfect duties are not:

We might be inclined to engage in—and even accept—moral excuse making of the following sort: ‘I know I should help the poor people in Africa, but I’m already stretched thin with all the money I am giving to the Red Cross’. But we wouldn’t be as inclined to make—or accept—excuses such as [this]: ...‘True, I shouldn’t have slapped that guy, but think of all the other people I didn’t slap’. The best explanation for this difference between what we regard as a valid excuse and what we do not, I suggest, is that it reflects our tendency to classify some duties as perfect and others as imperfect. We regard the duty to help those in need as imperfect and the dut[y] not ...to hit other people as perfect. (Stapleford, 2013: 4073)

As Stapleford notes, this sort of excuse-making seems perfectly acceptable in evaluations of rationality. Someone might excuse himself by protesting, “I didn’t deduce that particular theorem from the axioms, but look at all the other theorems I was able to prove!” or “I didn’t notice that there were jackdaws flying

¹⁷Nelson makes a similar response, though he phrases his argument in terms of act tokens and types on pp. 90-92 of the same article. Stapleford also uses the example of generosity on pp. 4073-74 of his article.

in front of me, but I noticed that there were birds flying in front of me. And anyways, I was trying to figure out whether we were close enough to hail that taxi.” Our ability to make excuses of this sort suggests that these rational norms are imperfect duties rather than perfect ones.

It is also worth noting that, just as we cannot predict which propositions among (1)-(4) a perfectly rational human would believe merely by knowing her evidence (that is, her epistemic circumstances), neither can we predict which particular actions a perfectly generous human would perform merely by knowing her moral circumstances. Although Nelson did not cast his argument in this light, it’s plausible that this inability to predict an ideal agent’s behavior, like Stapleford’s ability to be excused, is another mark of an imperfect duty.

Stapleford’s (2013) paper is directed against Nelson’s broader claim that we have no positive epistemic duties whatsoever. Nevertheless, his critique leaves intact Nelson’s more modest conclusion that there are no propositions such that we are obligated to believe *that particular proposition* on purely epistemic grounds.¹⁸ In fact, Stapleford’s analysis concerning the markers of imperfect duties gives Nelson a way to defend the modest conclusion against an apparent objection.

We can conclude then that there are at least some propositions that are permissible yet not obligatory to believe. Since (E-RUE) entails that it is rationally permissible to believe a proposition iff it is also rationally obligatory, (E-RUE) is false.

If (E-RUE) is false, as I have argued it is, this reveals an interesting way in which perfectly rational, evidential peers can arrive at different conclusions: they can disagree weakly. One can believe (or disbelieve, or withhold belief toward) a hypothesis whereas another can have no opinion concerning the hypothesis.

Rejecting (E-RUE), however, is not the same thing as endorsing rational permissivism. Building a case for permissivism requires rejecting (RUE) as well, which entails only that two perfectly rational, evidential peers *who have formed a judgment on a hypothesis* share the same doxastic attitude. More than that, arguing for *extreme* rational permissivism requires showing that evidential peers may rationally arrive at *contradictory* beliefs about a hypothesis. With these targets in mind, we turn to the problem of the best and the argument from abduction.

3 The Problem of the Best: The Falsity of (RUE)

3.1 Epistemic Bestism

Given that we’ve reflected on a proposition, what doxastic attitude ought we to have toward it? One plausible answer might be, “the best doxastic attitude.” We can call this “epistemic bestism.”

Epistemic Bestism For all S and for all p, if S has a doxastic attitude toward p then S ought to have the best doxastic attitude toward p given her evidence.¹⁹

¹⁸These seem like the sort of duties in which Nelson is most interested anyway.

¹⁹As with uniqueness, there are a number of candidate bestism principles that vary with respect to strength and range:

Expansive Epistemic Bestism For all S and for all p, S ought to have the best doxastic

The “ought” in each of these principles should be understood in the epistemic rather than moral sense.²⁰ Epistemic bestism might not seem very informative. If someone were to quote any of these principles to me in response to the question, “what ought I believe?” I should probably be annoyed: “If I knew the best doxastic attitude, I would not have asked!” Unfortunately, however, obnoxious advice is no less true for being obnoxious. Even if this principle does not count as informative advice, it can still be good encouragement when we do know what the best doxastic attitude is. If I am tempted to be cowardly, I benefit from the admonishment, “Be brave!” even if I know that is what I should do. When I am tempted as an epistemic agent to succumb to sentimentalism or am struggling to overcome implicit bias, I could benefit from the admonishment, “Remember, aim for the best doxastic attitude!”²¹

Despite the initial plausibility of epistemic bestism, I will argue in the following sections of the paper that it is mistaken. Often there isn’t any objectively best doxastic state or attitude to have toward a proposition. This is important because each of the uniqueness principles we have considered—when combined with a minimal number of plausible assumptions—entails a version of epistemic bestism.

Here is an argument for the conclusion that (RUE) entails epistemic bestism. Suppose that (RUE) is true and epistemic bestism is false. If (RUE) is true, but bestism is false, then there is some subject S such that for some proposition p, there is only one doxastic attitude that is permissible for S to bear toward p, and that attitude is not the best one. This could happen in one of two ways: either (i) a ranked attitude other than the best attitude toward p is alone rationally permissible, or (ii) an unranked attitude toward p is alone rationally permissible. On the first alternative, the second or third best (or fourth, or fifth, etc. . . .) doxastic attitude is alone rationally permitted. That the second or third best attitude is rationally permitted might not seem immediately problematic; however, since only one doxastic attitude is rationally permitted given (RUE), this means that adopting the best doxastic attitude toward a proposition is rationally forbidden. But this is surely absurd! (Compare: “Doing the morally best thing is morally forbidden.”)

The second way that (RUE) could be true and bestism false is if there simply is no ranking of the various doxastic attitudes (or at least among the ones at the top), and yet only one of these unranked attitudes is permitted. This option doesn’t wear its absurdity on its sleeve as obviously as the first alternative. It is, however, a strange position to maintain. For one thing, if only one doxastic attitude is permitted, we expect there to be a *reason* that particular attitude is privileged. But if there is no well-ordered ranking of the doxastic attitudes, then

state toward p given her evidence.

Strong Epistemic Bestism For all S and for all p, S ought to have the best doxastic attitude toward p given her evidence.

²⁰Clifford famously argues that such rules are not only epistemically but also morally normative. For the purposes of this essay, we will restrict ourselves to the purely epistemic interpretation.

²¹Someone might suggest that the best doxastic attitude toward any proposition just is to believe it if it is true and disbelieve it if it is false. Indeed, there might be an epistemic duty to believe truths and disbelieve falsehoods. This is not, however, the sort of duty anyone can have in mind when discussing rationality, for it is possible that a belief be both rational and false.

the uniquely permitted attitude is *no better* than some of the alternatives. And if no attitude is better than all the others—regardless of what criteria we use to determine what makes a given attitude better than another—there doesn’t seem to be a reason that *that* attitude alone is permissible. We could always ask why *that* attitude, rather than any one of the other unranked attitudes that are no worse than it, is the uniquely permitted one. Conversely, if some doxastic attitude alone is rationally permissible, we expect that there is something especially good *about that doxastic attitude*, given its evidential context, that lends it its privileged status. It shouldn’t merely be a brute fact about the world, nor should it be based on something other than facts about the privileged doxastic attitude or its relation to competing attitudes. On either alternative, the combination of (RUE) with epistemic bestism leads to unacceptable conclusions. Therefore if (RUE) is true so is epistemic bestism.²²

Epistemic bestism entails that, at least when S has a doxastic attitude toward a proposition, there is one that is best. Despite the initial plausibility of these theses, I will argue that epistemic bestism is ultimately untenable. Surprisingly, there is often no uniquely best doxastic attitude toward a given proposition that is uniquely fixed by the evidence. If there is no uniquely best attitude toward a proposition, it is impossible to have the best attitude toward that proposition. Bestism is false and so also rational uniqueness. The impossibility stems from the problem of maximizing over multiple, incommensurable variables. It is to this problem that we now turn.

3.2 Maximizing over Multiple Variables

Who is the biggest person in the world?

This is a tricky question. “Being big” may refer to mass, height, weight, width, volume, or any other size-related category. Which particular “bigness” is being assessed might often be clear in context, but the word itself invokes no distinction. Suppose Achilles is taller than Hercules, but Hercules is heavier. Who is bigger? It isn’t clear what the right answer is, or even that there *is* a right answer. The difficulty is that we do not know how much relative value to give height and weight. Neither height nor weight can be reduced to the other; neither can they both be reduced to some third quality. Height and weight are *incommensurable*. We cannot answer questions about superlative or comparative ascriptions about “bigness” (without contextual constraints) because one cannot maximize over multiple, incommensurable variables without an equivalence relation.

Is an equivalence relation anywhere to be found? Perhaps we could stipulate one. Supposing we had stipulated some equivalence relation, we could say whether Achilles or Hercules was bigger. For instance, if we posited that 1 inch equaled 5 pounds for the purpose of some metric, we could convert Achilles’ and Hercules’ height and weight into a single unit. The difficulty is that any stipulated equivalence relation seems arbitrary. Why shouldn’t 1 inch be worth 6 pounds, or 7 pounds? A dilemma emerges: (1) if we don’t stipulate an equivalence relation, we cannot make sense of the question, but (2) if we do stipulate an equivalence relation, the relation is arbitrary. No one can be the biggest person

²²The reasoning above can be replicated with minimal adjustments to show that if (E-RUE) is true, so is expanded epistemic bestism and that if (S-RUE) is true, so is strong epistemic bestism.

without a stipulated, arbitrary equivalence relation. More generally, without such an equivalence relation, nothing does or even could satisfy the superlative of an incommensurable, multi-variable expression.

At any rate, *almost* nothing could. Suppose there were someone—let’s call him Goliath—who was bigger than every other person with respect to every single size-related category. It might truly be said of Goliath that he is the absolutely biggest person. Or at least, it can truly be said of Goliath that he would qualify as the biggest person on any stipulated equivalence relation, and this seems like a good enough reason to *call* him the biggest. Either nothing can satisfy the superlative of an incommensurable, multi-variable expression or nothing can satisfy the superlative of an incommensurable multi-variable expression unless it satisfies the superlative of every variable within the multi-variable expression.

There is not very much at stake over whether Achilles or Hercules is bigger. But the problem of maximizing over multiple variables is neither new nor without importance to philosophy. One typical application is as a reason to reject an early version of utilitarianism according to which an action is right iff it maximizes pleasure and minimizes pain. Critics have pointed out that it is impossible in some situations both to maximize pleasure and to minimize pain since neither is reducible to the other. (Minimization is no less difficult than maximization in this regard. Minimization is merely maximization of an inverse.) Moreover, utilitarians have for the most part tried to accommodate this critique, finding ways to evaluate utility by a single variable (preference, for example).

So far, we have been considering the word, “biggest.” Identifying the biggest is difficult because “being big” can indicate height, weight, mass, or a host of other qualities. Another word that often operates similarly is “best.”

Let’s modify an example from Plato’s *Meno*. In Plato’s famous dialogue, Socrates asks Meno what value there is in knowing the road to Larisa over and above having the right opinion about the road to Larisa. This is a difficult question, requiring us to distinguish knowledge from merely true belief. But Socrates’ question would have been even more difficult if he had asked about the best road to Larisa, for in order to answer this question, we need not just understand what knowledge is, but how to determine which road is best.

In order to understand what is the best road, one must know what the good-making properties for roads are. The difficulty of determining one, unique, best road, is that there seem to be several (potentially) conflicting and incommensurable good-making properties for roads. Let us consider just two of these qualities: being quick and being easy. Let us suppose that road A is quick but not easy, road B is easy but not quick, and a third road, road C, is quicker than B but slower than A and easier than A but more difficult than B. Which road is best? There seems to be no answer to this question.²³

²³Questions about maximums can be difficult even when multiple variables are not involved. Suppose Opie is 4’6”, and both Andy and Barney are 6’2”. Who is tallest? There is no good answer to this question either. Or rather, there is a good answer to this question, but it is bad English. No one person is uniquely the tallest, although Andy and Barney are both tied for being the most tall. The difficulty of this question, however, is fundamentally different than the difficulty about the best road. In the question about height, we know that Andy and Barney are tied because they have the same height. The only difficulty is that the use of a superlative in English prompts us to pick out one unique object, and in this case there is a tie between two. In the question about which road is best, we don’t even know

Note that this is a different question than “Which road is best for so-and-so?” Perhaps Socrates needs to get to Larisa very quickly, and Meno’s leg is cramped. Even if there is no answer to the question of which road is best *simpliciter*, we are still prepared to say that the best road for Socrates is road A and for Meno B. Why is this so? One plausible answer is that it is because Socrates and Meno have different telic or pragmatic values assigned to the various good-making properties of roads. In this case, Socrates values the goodness of speed over ease whereas Meno is presently disposed to value ease over speed. The values that they have actually assigned (or their preferences, or their dispositions) give us a way to weigh the different values, but they only do so relative to an individual’s values and goals, not absolutely.

3.3 The Argument from Abduction

Abduction may be casually defined as “inference to the best explanation.” This definition is too simple. Various restrictions must be put in place. If hypothesis H1 is the best overall, but H2 is only slightly worse than H1, then an abductive inference to H1 is invalid. So not only must H1 be better than all other (considered) explanations, but it must also be significantly better than the second best explanation.²⁴ A related constraint often suggested is that the hypothesis must be all things considered probable in the eyes of the agent. If H1 is the best explanation and significantly better than its competitors, but improbable overall, then an abductive inference to H1 is invalid. Although it is important to acknowledge these constraints, abduction nevertheless begins with a search for the best explanation.

How does one determine which theory is best? No canonical list exists, but there is general agreement on some of the core considerations. All else being equal, a theory is better if it...

1. Is simple
2. Fits the data
3. Explains the data
4. Is elegant

Let us call these qualities “theoretic virtues” and examine each further.

Simplicity The value of simplicity in philosophical thinking is most famously captured by Ockham’s razor. More recently, we can call to mind Quine’s articulated preference for desert landscapes. All things being equal, we prefer a theory that does not posit many things when positing a few would suffice. We call this kind of simplicity “parsimony.” Simplicity itself can be broken into distinct desiderata. Simplicity of types is distinct from simplicity of tokens. David Lewis’s version of modal realism, for instance, has type but not token simplicity.

whether or not there is a tie. Answering that the three roads are equally good would be no better than answering that one is best indisputably. Any answer we give would imply the commensurability of good-making qualities that are, in fact, incommensurable.

²⁴Just how significant the gap must be is also a matter of debate.

Fitting the Data Especially when dealing with theories that attempt to account for a wide spectrum of phenomena (as in scientific theories), we often find that no available theory perfectly fits the data. Nevertheless, we value theories that have few outliers and anomalies.

Explaining the Data Explaining the data is distinct from fitting the data. Imagine trying develop a theory for some evidence set $\{E_1, E_2, E_3, \dots E_n\}$, where all the evidence is propositional. One theory that would fit this data very well would be the theory that of all the propositions that are members the set are true. The theory postulates that the truth of each proposition as brute fact. However appalling we might find this theory, we must admit that it would fit the evidence perfectly. Nevertheless, such a theory would not explain any of the data but only restate it. We value theories that explain much of the data by relying on relatively few brute facts and primitive concepts.

Elegance Simplicity, no doubt, contributes to a theory's elegance, but we also value other aspects of a theory that make it lovely. We value theories that have patterns, symmetry, or parallels, and theories that are unified rather than disjointed. It is perhaps not clear why we epistemically favor these theories, but that we favor them cannot be doubted.

Some philosophers would prefer to do away with talk of inference to the best explanation and talk instead about inference to the most probable theory. But it is unclear how probability, subjectively understood, could be determined without relying on appeals to the theoretic virtues listed above, or others like them. If asked why a certain theory seems more probable, it seems natural to respond by saying that it is simpler, or better fits the evidence, or is more elegant, etc. Thus even if inference to the best explanation is ultimately cashed out in terms of probabilities, the same cluster of theoretic virtues are likely to play an important role in our epistemic decisions.

Abduction plays an important role in our ordinary reasoning. This morning in class, a student who is honest and has excellent attendance was not present. I also had an email in my inbox from the same student saying that he was sick and could not attend, and he had shown signs of having a cold during our previous class. One relatively simple explanation is that he was indeed sick and sent me an email to confirm his absence. The evidence was not, however, incompatible with other hypotheses. It could have been, for instance, that he did not come because he was kidnapped by friends for a day in the city, and that his friends also emailed me from his account to let me know that he was "sick." The cold the previous day could have been a fortuitous (for the sake of confirming his story) coincidence. Both hypotheses are compatible with the evidence, but it is much simpler to suppose that he actually was sick than that his friends hatched an elaborate plot that happened to coincide with his earlier symptoms. There is an abductive reason to prefer his story to the more elaborate explanation.

Abductive reasoning is also essential to anti-skeptical arguments. Few philosophers believe the existence of the external world can be proven deductively, and Hume's problem of induction is a serious challenge to attempts to do so inductively. A plausible strategy is to show that the existence of the external world offers the best explanation for our apparently coherent and overlapping experiences. I do not claim that abduction solves the problem of skepticism, but it

does seem to give the philosopher more ammunition against certain skeptical problems.

Finally, abductive reasoning is integral to science. Consider the following example from Igor Douven:

At the beginning of the nineteenth century, it was discovered that the orbit of Uranus, one of the seven planets known at the time, departed from the orbit as predicted on the basis of Isaac Newton's theory of universal gravitation and the auxiliary assumption that there were no further planets in the solar system. One possible explanation was, of course, that Newton's theory is false. Given its great empirical successes for (then) more than two centuries, that did not appear to be a very good explanation. Two astronomers, John Couch Adams and Urbain Leverrier, instead suggested (independently of each other but almost simultaneously) that there was an eighth, as yet undiscovered planet in the solar system; that, they thought, provided the best explanation of Uranus' deviating orbit. Not much later, this planet, which is now known as "Neptune," was discovered. (Douven, 2011)

These examples demonstrate what Igor Douven calls the "ubiquity of abduction" (Douven, 2011). It is present in ordinary, philosophical, and scientific reasoning. If the nature of abduction creates problems for rational uniqueness, those problems will be ubiquitous as well.

Indeed, abduction does create a problem for (RUE). Abduction relies on determining a best explanation. But as we have seen, there are several incommensurable criteria for determining which theory is "best." Our previous argument has shown that maximizing over multiple variables is impossible. Accordingly we cannot speak of an absolute best, but only what is best relative to a certain weighting of the theoretic virtues. The distinct theoretic virtues will sometimes be competitive, and as such trade-offs between virtues will be necessary. Someone who values simplicity especially highly may come to different conclusions than someone who values more highly the strict fitting of the data.

Unlike the argument from efficient evidence, the argument from abduction results in the rational permissibility of strong disagreements among evidential peers. Suppose Amy endorses H1 on the basis of evidence E according to her weighting of the theoretic virtues whereas Nico endorses the incompatible H2 on the basis of the same evidence E and his distinct weighting of the theoretic virtues. (Perhaps Nico values the token simplicity of a theory more than does Amy.) Amy is committed to the denial of H2 just as Nico is committed to the denial of H1. There will be propositions that Amy rationally believes such that Nico believes the negation. Yet both may do so rationally if they have appropriately evaluated the hypotheses according to their distinct weighting of theoretic virtues.

The ubiquity of abduction makes it an indispensable part of our reasoning practice. Accordingly, we cannot reject abductive inference because no best explanation exists in the absolute and universal sense that (RUE) requires. Rather, we must interpret the inference to the best explanation as an inference to the best explanation according to an individual's standards, that is her weighting of the theoretic virtues. Having considered the argument from efficient evidence and the argument from abduction, we have good reason to believe that there

can be rational disagreement among evidential peers. Whether a disagreement is committed or incidental, strong or weak, an agent's interests and theoretic values can lead to diverse and contradictory doxastic attitudes without either party having made an error of reasoning.²⁵

3.4 Is Extreme Permissivism Too Permissive?: Defending the Argument from Abduction

The theory I have sketched is permissive indeed. Some might think that it is a good deal too permissive, and that I have opened the floodgates to a multitude of unacceptable doxastic attitudes.

Says the objector:

On your view, it seems that almost *any* view could be rationally permissible. At least, any view can be permissible so long as one has the right sort of interests, cognitive goals, and epistemic values to go along with it. It could be rationally permissible for someone to be a solipsist, or to doubt the existence of an external world. But it should be obvious that some views are just plain unacceptable. If your view has the consequence that a view as bizarre as solipsism could be permissible, so much the worse for your theory! Your permissivism is simply too permissive.

My view really does have the consequence that all sorts of beliefs might be permissible given the right epistemic framework. After all, despite its strangeness, solipsism is a very parsimonious theory. It posits only one object, my own mind! And what if someone were to value parsimony far enough beyond the other theoretic virtues that solipsism started to look like a good theory—perhaps even the *best* theory? Such a solipsist would be rational on my view. Why isn't this a crazy thing to say?

The first thing to note is that not every epistemic failure must also be a failure of rationality. As Kvanvig rightly warns us, “there is a predilection among epistemologists to identify too many regrettable features of cognitive systems in terms of one favorite term, ‘irrational’” (Kvanvig, 2013: 235). Solipsism has the epistemic defect of being false. Worse than that, it upholds a way of looking at the world that ensures one will come to many more false beliefs. There is a sense in which solipsism is thus deeply delusory. But we should hesitate to jump immediately from delusion to irrationality. A system of beliefs may be delusory in virtue of dramatically clashing with the way the world is and not necessarily because the agent has been epistemically irresponsible. Systematically irrational epistemic behavior may be a surefire way to end up with delusory beliefs, but the reverse may not always be true.

Consider for instance the false yet rational beliefs we would hold if we lived in Descartes' demon world. Our beliefs would be epistemically deficient in many ways. They would be false. Our belief-forming processes would be unreliable. Our belief system would be delusory. And yet our beliefs would continue to be rational.

²⁵Though developed separately, this argument has some parallels to an argument from Igor Douven's (2009) “Uniqueness Revisited.” *American Philosophical Quarterly* 46: 347–361. Here I try to place the argument within a larger, permissivist strategy, and, in what follows, respond to uniqueness-friendly rejoinders.

If a solipsist's beliefs really are calculated correctly given her cognitive goals and epistemic values, they are rational. Of course they are profoundly wrong and delusory. But it is not uncommon to form false, rational beliefs. Why should it be surprising that solipsism is among the false beliefs one can rationally believe, given the right sort of epistemic framework? Nearly no one (perhaps no one at all) actually has the sort of epistemic framework according to which solipsism could be permissible, but it *would* be permissible for someone who *did* have the that framework.

Replies the objector:

Your "solution" only pushes the problem back a step. It seemed crazy that solipsism could be rational. You responded that it could be rational if someone valued a certain kind of parsimony far and above the other theoretic virtues. But this weighting of the theoretic virtues is surely just as crazy! If someone really values parsimony that much, to the extent that he doesn't believe anything outside his mind exists, he is behaving irrationally. That is an incorrect way to weigh the theoretic virtues.

If there are incorrect ways to weigh the theoretic virtues, there must also be a correct way to do so. The objector insists—and it is a tempting line of thought—that there *must* be a fact of the matter about how parsimony and explanatory power, for instance, ought to be weighed against each other in order to avoid the rational permissibility of idiosyncratic weightings. That there are such facts is just what this paper denies.

The objector presses: "Why think that there is no fact of the matter about the best way to weigh the theoretic virtues?" Here is a simple argument:

1. If the theoretic virtues are incommensurable with each other, then there is no fact of the matter about the best way to order them.
2. The theoretic virtues are incommensurable with each other.
3. Therefore, there is no fact of the matter about the best way to order the theoretic virtues.

This paper has already said a great deal in defense of (1) in section 3.2. It is impossible to maximize over multiple, incommensurable variables without a purely stipulative, equivalence relation. We noted a possible exception for "Goliath" cases, but we are rarely if ever in situations in which one doxastic attitude is better than its competitors with respect to *every* rational good-making quality. And it would be very odd if, although there is no objectively best way to order the theoretic virtues, the norms of rationality required a precise ordering. In the next section, I shall consider the possibility that objective rules of trump can produce a uniquely correct ordering of the theoretic virtues even if the theoretic virtues are incommensurable. Temporarily bracketing this objection, however, (1) looks like a strong premise. The uniqueness's best bet is to challenge premise (2).

So why think that the theoretic virtues really are incommensurable? One reason, albeit a relatively weak one, is simply our inability to know how such virtues commensurate. One theory that explains our ignorance is the thesis that there just is no fact of the matter about how such values commensurate to be

known. Of course other explanations are possible. I am not sure how many farthings a shilling was worth in old English currency, but I have no doubt they were commensurable. A stronger reason is that questions about the worth of one theoretic virtue in terms of another seem not just difficult but nonsensical. “How many farthings are a shilling worth?” may be a confusing question for those of us born after the decimation of the British currency, but “how many inches is a pound worth?” sounds like nonsense. Farthings and shillings are both units of the same kind of thing, namely English currency. One can (or rather could) say of someone’s income that it is worth x number of farthings or that it is worth y number of shillings without having said anything different. Both claims would be equivalent. Inches and pounds, however, are measures of completely different sorts. One measures length, the other weight. There are no values for x and y such that the sentence that a box measures x inches in some direction and the sentence that the same box measures y pounds express equivalent truths. Units of simplicity (“simplicitons”) and units of explanatory power (“powerandums”), like inches and pounds, measure completely different kinds of qualities. The same theory may be evaluable both in terms of its measure of simplicity and its measure of explanatory power, just as the same box may be evaluable both in terms of pounds and inches. But assertions about the simplicity of a theory are never equivalent to assertions about its explanatory power.²⁶ This gives us good reason to think that simplicity and explanatory power really are incommensurable, and consequently, that there is no objectively correct way to weigh them against each other.

Are there then *no* rational constraints at all on the ways one can weigh various epistemic goods? Even the incommensurability of epistemic goods does not entail such a strong thesis. Consider the following principles:

(P1) In order to be rational, one must value every theoretic virtue.

If someone doesn’t value simplicity as a virtue for theories, she isn’t being rational.

(P2) In order to be rational, one must not value any theoretic vice.

If someone values acquiring falsehoods as such, or counts it as an epistemic good for a theory that he wants it to be true, he isn’t being rational. Furthermore, even given that all and only the epistemic goods are valued, not just any hypothesis could be rationally permissible according to some weighting of values:

(P3) If doxastic attitude A concerning p is worse than doxastic attitude B concerning p with respect to one epistemic good, and is worse than or equal to B with respect to all other epistemic goods, then doxastic attitude A concerning p is impermissible.

Consider the theory that I haven’t received my coffee yet because the barista has forgotten my order due to the whimsy of an invisible, meddling, memory-erasing wizard. This theory is less simple than the theory that I haven’t received my coffee solely because the barista has forgotten my order, and it enjoys no other epistemic advantages over the simpler theory. Accepting the wizard hypothesis

²⁶At least, they are never equivalent when those assertions have contingent content.

is thus rationally impermissible according to any way of weighing the epistemic goods.

(P1)-(P3) are not meant to be exhaustive. They are simply meant to show that accepting the incommensurability of epistemic goods does not entail the absence of any rational constraints on permissible ways of weighing epistemic goods or of rationally permissible hypotheses.

Extreme permissivism is permissive indeed, but not recklessly so. There will be some bizarre hypotheses that, when conjoined with a certain way of weighing epistemic goods, cannot be denounced as irrational; however, this seems less problematic when we consider that they can be epistemically criticized in other ways. Moreover, the extreme permissivists can still place substantive constraints on the acceptable weightings of epistemic goods without violating their intuitive incommensurability.

3.5 Tricky Trumping

There is at least one strategy left to the defender of (RUE) who takes the incommensurability of various epistemic goods seriously. The final objection we shall consider concerns the possibility of trumping.

The card game of spades provides an excellent example of how trumping works.²⁷ A round of spades is composed of 13 “tricks,” during which all four players successively play a card. The highest card played wins the trick, and then whoever wins that trick plays the first card of the next trick. In spades, suit matters. The cards of different suits are incommensurable in value. Though a seven of hearts is higher than a 6 of hearts and lower than an eight of hearts, it is neither higher than nor lower than nor equal to a five of diamonds. But this creates a problem; for spades is only playable if something we may call “Trick Uniqueness” is true:

Trick Uniqueness For any ordered set of cards OS and for all cards C, if OS is played, there is just one card C such that C wins the trick.

If Trick Uniqueness were false, there would be times when nobody could win the trick. In fact, as it is not uncommon for cards of different suits to be played in the same trick, many of the tricks would not be winnable. The game would stutter to a halt any time two different suits were played.

The game does not stutter to a halt because in spades trick uniqueness is true. But how can this be? For how can there be a highest card if the different suits are incommensurable? And hasn't our argument hitherto assumed that this is impossible?

The problem is solved by introducing rules of trump. As the game's name suggests, spades is the trump suit, so that any spade beats any card of any other suit. Each spade can only be beaten by a higher spade. So for instance, in the ordered set of cards {QH, 5S, 10D, 3S}, the five of spades wins. It trumps the queen of hearts and the ten of diamonds, and it is higher than the three of spades. It is not that spades are now commensurable with other cards. A two of diamonds is *one* lower than a three of diamonds, and *nine* lower than a Jack of diamonds. But even though we know a two of spades beats a two of diamonds, there isn't a good answer to the question, “By how much?” Spades

²⁷And it is an excellent game to boot!

are on a completely different level than diamonds. This is what it means for spades to *trump* them.

Yet a problem remains: sometimes a trick contains multiple suits but no spades. To fix this problem, the game's inventors introduced another rule of trump. The suit of the first card played in a trick trumps all suits other than spades. Thus in the trick {JH, 6D, 3H, QC}, the jack of hearts wins, whereas if the order had been slightly altered to be {6D, JH, 3H, QC}, the six of diamonds would have won.

If rules of trump can be introduced to secure trick uniqueness despite the incommensurability of different suits, perhaps analogous rules of trump can be introduced to secure rational uniqueness despite the incommensurability of different good-making properties for doxastic attitudes.

It is difficult to rule out this alternative completely. In part, this is because there are an endless number of candidate rules of trump that (RUE) advocates can present for consideration. The task of ruling out all of them may turn out to be equally endless. Nevertheless, using spades as a foil, let us examine some of the more promising trumping strategies on behalf of (RUE).

Let's consider an example. Suppose that when evaluating a theory, the epistemic virtues of simplicity and explanatory power are in competition with each other. Does one of these goods trump the other? It seems that neither one does. In a game of spades, a two of spades always beats an ace of diamonds. It doesn't matter that the two is the lowest spade or that the ace is the highest diamond. The spade is trump; it beats all comers. If simplicity always trumped explanatory power, it would never be rational to believe in any complicated theory. If explanatory power always trumped simplicity, it would never be rational to avoid explaining outliers and other anomalous data, even if doing so meant achieving significant theoretical simplicity. But these conclusions are absurd. Neither simplicity nor explanatory power outright trumps the other. And the same could be said with respect to any of the other epistemic goods. No single theoretic virtue trumps all comers.

This flat-footed strategy will not work for (RUE), but perhaps a more nuanced version will. According to spades, any spade trumps all other non-spades. But the rules need not have been so. In a similar game, a game we may call spades*, a card of spades is only trump if it meets a certain threshold. Perhaps only the 10 of spades and higher trump.

First, it is unclear that this strategy will be able to save (RUE). In spades*, if both the rule of trump regarding spades and the rule of trump regarding the first suit are "threshold trumps" rather than "outright trumps," Trick Uniqueness may turn out to be false. In spades, the rules of outright trump guarantee that every trick of mixed suits has a rule of trump to govern it. Not all rules of threshold trump do the same. If the threshold for spades-trump was a ten of spades, and the threshold for first-suit-trump was a five of the first suit, the following trick would have no outright winner: {4D, 9S, 10H, 8H}. If spades* is to be a playable game, there must be some additional rule that makes Trick Uniqueness true. Similarly, if this strategy is to be successful for (RUE), there must be some additional explanation for why rational uniqueness turns out to be true.

Second, the threshold strategy does little better than its flat-footed counterpart in light of our actual reasoning practices. Is there a level of simplicity such that any theory that meets this threshold trumps all theories that fall short of

it? Or of explanatory power or of elegance?

Imagine that there are some hypothesis H_1 and H_2 such that H_1 exactly meets the trumping threshold of theoretical virtue TV_1 . H_2 is minimally worse than H_1 with respect to TV_1 but is significantly better than H_1 with respect to TV_2 . Given the stipulated trumping threshold, H_1 must be ranked higher than H_2 , no matter how small the difference between the hypotheses with respect to TV_1 , nor how great the difference with respect to TV_2 . This might mean, for instance, that there would be some cases in which it would be irrational to believe a theory that is slightly less simple, but massively more explanatorily powerful than its best competitor. At the very least, this consequence is counter-intuitive. Which theoretic virtue shall we substitute for TV_1 such that this story seems plausible when every other theoretic virtue is substituted for TV_2 ? Yet this is the threshold-trumping advocate of (RUE) must say.

Perhaps the defender of (RUE) will make the following counter:

You've been assuming that whatever quality trumps must always be the trumping quality. But there is an alternative. Perhaps the trumping that preserves (RUE) is less like spade-trumping and more like first-suit-trumping. It isn't that there is one quality such that, when a certain threshold is met, it trumps all comers. Rather, various contexts, orderings, or epistemic circumstances select different qualities and corresponding thresholds that act as trump *for that context*.

This speech should not move us. The same difficulties re-emerge in only slightly different guise. First, it is unclear whether this maneuver genuinely preserves uniqueness. According to (RUE), *only* evidential factors determine what is rational. But, assuming that the context is something beyond the evidence itself, this maneuver cannot preserve rational uniqueness. For (RUE) to be true, what is rationally permissible cannot be fixed the evidence *and also* facts about the context. If the various contexts, orderings, and circumstances are non-evidential in character—and it is hard to imagine how they could not be—then far from providing a defense for (RUE), they will provide an explanation for why (RUE) is false. Second, it is not clear that this strategy would be an effective escape from the charge against context-independent, threshold trumping, even if the relevant context were somehow solely determined by the evidence. Imagine there is some hypothesis H_1 and some hypothesis H_2 such that H_1 exactly meets the trumping threshold of theoretical virtue TV_1 for context C. Again, there will be situations in which H_1 must be ranked higher than H_2 , no matter how small the difference between the hypotheses with respect to TV_1 , nor how great the difference with respect to TV_2 . Perhaps in a different context, the trumping threshold for TV_1 would be different. Or perhaps some other theoretic virtue would trump it. But this should not change our intuitions about the particular case. The fundamental problem in both formulations is that dramatic differences with respect to a theoretic virtue should be irrelevant to considerations about whether or not a belief is rationally permissible.

It would be painstaking indeed to exhaust the potential rules of trump a defender of (RUE) might propose for consideration. Hopefully, however, we have cast doubt on the more promising strategies. The (RUE) advocate is invited to try her hand at a game of trump; but, the more complex and tricky

the rule of trump, the less likely we are to see it as a simple, elegant, and feasible theory of rationality.

3.6 Just More Evidence?

This paper has so far been silent on the nature of evidence. It hasn't weighed in on such questions as whether evidence is ultimately phenomenal or whether our evidence is what we know. Nor will this paper wade into such questions more deeply now. No matter what evidence is, there must be some way that evidence is evaluated. If this paper is right, then differences in the way evidence is evaluated (e.g. different cognitive interests or different epistemic weightings of the theoretic virtues) can result in different rational attitudes even when there is no difference in evidence.

The paper's silence on the nature of evidence has, however, left it open to a dangerous question: what if the cognitive interests and epistemic values themselves are part of the evidence? ²⁸

Can we identify a difference between evidence on the one hand and cognitive interests and epistemic weightings on the other that can safeguard this project against such questions?

It should be fairly clear that cognitive interests, at any rate, are not themselves part of the evidence. Evidence can rationally affect the confidence one is permitted to have in a hypothesis. But in normal circumstances, an agent's interest in some proposition does nothing to influence how much confidence it is rational to invest in a hypothesis. (There may be abnormal cases such as when the proposition in question is about the agent's interests. But we need not trouble ourselves overly much about these: the normal cases are enough to secure the intended conclusion.)

Imagine for instance that Ahmed is weighing the evidence bearing on the question of whether or not it will rain. Evidence that it will rain and Ahmed's interest in the question of whether it will rain both make it more probable that Ahmed will believe that it will rain. But they accomplish this in entirely different ways. Ahmed's having evidence that it will rain increases the chance that Ahmed believes that it will rain and decreases the chance that he believes that it will not rain (assuming that Ahmed is behaving rationally). Ahmed's interest in the question of whether or not it will rain makes it more likely that he will consider the evidence bearing on the question of whether or not it will rain, but such interest by itself rationally favors neither the rain nor the non-rain hypothesis for Ahmed.

Epistemic weightings of the theoretic virtues present a trickier case. Clearly, the amount of evidential support Ahmed has for the proposition that it will rain is unaffected by his interest in the question. But the credence one is licensed to invest in a hypothesis might well be affected by the relative value one gives the theoretic virtues. Simple hypotheses should be considered more probable by those who value simplicity highly, for example.

Cases in which an agent is unreflective about her standards, however, support the thesis that the epistemic values by which one weighs the evidence are not part of the evidence themselves. Let's suppose again that Ahmed is weighing the evidence for and against the proposition that it will rain. Ahmed weighs

²⁸I owe thanks to an anonymous reviewer for pressing this point.

his evidence according to the way that he weighs the various theoretic virtues; however, Ahmed has no beliefs about the way that he weighs the various theoretic virtues. He does, of course, have some way of evaluating his evidence; it's just that he is unreflective about what his method of evaluation is.²⁹ According to many theories of evidence, evidence is or is at least partially constituted by either propositions or attitudes that contain or refer to propositions (belief, rational belief, knowledge, etc.). If any of these views is right, then Ahmed's weighting of the theoretic virtues is not a part of his evidence—he doesn't even have any beliefs concerning them!³⁰ But they do affect the way that Ahmed rationally weighs his evidence and so also the doxastic attitudes that are permissible for him.

A uniuquist might try to defend her view by suggesting a more expansive view of evidence. Perhaps one's evidence need not map perfectly onto one's occurrent mental states (whether those be knowledge, beliefs, etc.) but the mental states one would have if prompted under the right circumstances. (For instance, many of our background beliefs seem to play a role in rationally justifying our doxastic attitudes even though our background beliefs are often not presently occurrent when we form beliefs on the basis of them.)³¹ Then perhaps even if Ahmed does not actually have any occurrent beliefs (or knowledge, etc.) about the way he weighs the theoretic virtues, he would have accurate beliefs about them if prompted in the right sort of way.

This response leans heavily on what “the right sort of way” is, but in any case I think we can develop versions of Ahmed's story that circumvent such maneuvers. There may be psychological reasons that Ahmed is unable to acquire beliefs about the way he weighs the theoretic virtues. Perhaps Ahmed is still a child and has not yet developed the level of reflection necessary for answering or even understanding such questions. But even children can have rational beliefs! Or Ahmed may be disposed to have opinions about his own epistemic preferences that are systematically incorrect. Perhaps Ahmed was dazzled by the presentation of Ockham's razor in his introductory philosophy course and so is disposed to believe that he values simplicity above all other theoretic virtues; however, if one were to look at the theories Ahmed actually endorses, one would notice that he actually values simplicity less and explanatory power more than the average epistemic agent. These observations highlight the point that what attitude it is rational for Ahmed to have depends on Ahmed's epistemic values themselves, not on his mental states about his values. If one has a theory of evidence according to which evidence is partially constituted by mental states with propositional content, one should not think that one's way of weighing the theoretic virtues is a part of one's evidence.

But what if one's conception of evidence is expansive enough to include non-content-based mental states—expansive enough to include an agent's epistemic values or preferences, even when those states are never incorporated into a propositional attitude by the agent?

One reason for endorsing a broad conception of evidence like this might be a core commitment to evidentialism. According to evidentialism, the evidence that one has—and the evidence alone—fixes what it is rational to believe. Does

²⁹Most of us are often in Ahmed's position, I imagine.

³⁰There are, of course, still propositions about Ahmed's weighing mechanism, but they are not propositions to which Ahmed himself seems to have the right kind of access.

³¹For an excellent discussion of these and related issues, see Feldman (2004).

the way one weighs the theoretic virtues play a role in fixing what it is rational to believe? Then (says the evidentialist) the way one weighs the theoretic virtues must be a part of the evidence!

How should the permissivist understand this move? One unsympathetic interpretation of the maneuver is that it preserves evidentialism (and so takes strides toward preserving rational uniqueness) only by trivializing an otherwise substantive thesis about rationality into a definition of evidence. Since no matter what turned out to have an impact on rationally acceptable attitudes would be considered evidence, it would be impossible to develop a counterexample to the thesis that evidence alone fixes what is rational.

But this would be uncharitable. First, an evidentialist might have principled reasons for endorsing a broad conception of evidence. But more interestingly, it's at least conceivable that the best theory of evidence identifies evidence by its functional role in reasoning (e.g. by whether or not it fixes what is rational) as opposed to the kind of mental state that constitutes it (e.g. belief, knowledge, etc.).

One way to respond to such a view is to point out the distinct functional role that weighing mechanisms play in the reasoning process. Epistemic agents reflect on the evidence. But they reflect on the evidence according to the way they weigh the theoretic virtues. The relationship between evidence and the way one values the theoretic virtues is like the relationship between objects that are weighed and the scale by which they are weighed. Both work together to predict a weight, but their contributions to the process are sharply distinguishable. The object that is weighed provides the mass whereby force is applied. The scale doesn't do that. But the scale does register the force exerted by the object. Without both, one could not predict the weight of the object. Similarly, evidence is what confirms or disconfirms hypotheses. The theoretic virtues don't do that. But they are (at least part of) the criteria whereby the evidence confirms or disconfirms a hypothesis. Without both, one could not rationally evaluate a doxastic attitude.

Occasionally, facts about how an agent evaluates her evidence with respect to the theoretic virtues will themselves be evidence for a hypothesis. For instance, the fact that Ahmed values explanatory power more than simplicity, together with the fact that Ahmed values simplicity more than elegance, is evidence for the hypothesis that Ahmed values explanatory power more than elegance. So the fact that Ahmed values explanatory power more than simplicity can function both as evidence and as a description of the criteria by which Ahmed evaluates his evidence. Does this overlap of functional roles refute the thesis that such functions are sharply distinguishable?

No, that one can treat facts about an agent's criteria of evidence as evidence shows only that some facts may be capable of performing two distinct roles. Neither does the observation that scales are among the things one can weigh with a scale refute the thesis that the functional roles of having mass and registering force are distinct.

Admittedly, this metaphor is more suggestive than definitive. Saying more would likely involve defending a full theory of evidence, a task which is beyond the scope of this paper. There are, however, at least *prima facie* reasons for thinking that the way one weighs the theoretic virtues should not generally be considered a part of one's evidence, for such weightings play a functional role—that of being the criteria for evaluating the evidence—that is distinguish-

able from the functional role of evidence itself.

This section began by considering the objection that an agent’s cognitive interests and epistemic values might themselves be a part of an agent’s evidence. If this were the case, then noting that agents with different interests or ways of weighing the theoretic virtues can rationally endorse different hypotheses would do nothing to refute (RUE). Admittedly, this paper has neither considered every plausible theory of evidence nor offered one of its own, and perhaps for this reason, the conclusions of this project depend on further investigation into the nature of evidence. We did note, however, that cognitive interests cannot be a part of the evidence since they do nothing to affect the confidence one is rationally permitted to invest in a hypothesis. Moreover, if evidence is constituted by mental states that involve propositions, the evidence cannot include an individual’s weighting of the theoretic virtues, for they can affect what is rationally permissible for an agent to believe even when she does not (or cannot) have any mental states about them. Finally, even on a particularly expansive conception of evidence, there is reason to think of an individual’s ways of weighing the theoretic virtues as distinct from the evidence because of the different functional roles they play in epistemic decision-making.

3.7 Conclusion

In the absence of further argument, the prospects for rational uniqueness look grim. Over the last several pages, we have considered several reasons to reject both (RUE) and (E-RUE). The vast number of propositions justified by our evidence and the imperfect nature of positive epistemic duties give us reason to think it is at least sometimes permissible not to acquire every doxastic attitude that is permissible for us. This has led us to deny (E-RUE). We noted that if (RUE) is true, epistemic bestism must be as well. But since there are several distinct good-making properties of doxastic attitudes, and since one cannot maximize over multiple variables,³² neither can there be an absolutely best doxastic attitude given only the evidence. Given different weightings of the theoretic virtues, evidential peers might even develop strong and committed cases of rational disagreement.

In the next section I consider arguments for Uniqueness as presented by Roger White and Richard Feldman. Although their arguments force the permissivist to draw some important distinctions, they do not compel the permissivist to abandon her position.

4 Resisting Uniqueness

4.1 Quick and Dirty Permissivism

In “Epistemic Permissiveness,” White begins his defense of uniqueness with what he calls the “quick-and-dirty” argument (White, 2005: 447). Though we shall also evaluate some of his “long-and-clean” argumentation, this argument is a good place to begin as it captures one of the central intuitions motivating rational uniqueness. White asks us to imagine that we are in the juror’s box weighing

³²That is, except in “Goliath” cases, which do not appear to be at play in the argument from abduction.

the evidence for and against Smith, the defendant:

[T]he 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. (White, 2005: 447)

This is a natural and intuitive thought. In his (2013) paper, White expresses this as the idea that evidence is unidirectional. And, indeed, how could evidence be both for and against a hypothesis? How could evidence make something more and less likely at the same time?

Fortunately, if the arguments for permissivism previously presented are sound, then there is a quick-and-dirty response to White's objection. Evidence cannot be evaluated apart from *standards*. What exactly are epistemic standards? Epistemic standards, in the context of this paper, are the criteria according to which the support relations between evidence and hypotheses are evaluated. Such standards include the ways that an agent may weigh epistemic goods, including at least the theoretic virtues. (I see no reason to assume that facts about an agent's valuation of theoretic virtues must be exhaustive of what standards encompass. In Kelly [2013], for instance, Tom Kelly argues that what it is rational to believe depends in part on the way an agent weighs different cognitive goals, e.g. to believe truths or to avoid false beliefs.) The permissivist's quick-and-dirty response is this: no set of evidence can be evidence for and against *p* according to the same standards; a set of evidence can, however, be evidence for *p* according to one set of standards, but evidence for not-*p* according to alternative standards.

4.2 Feldman's Principle and Escalating Skepticism

In "Reasonable Religious Disagreements," Feldman anticipates this dialectical move from the permissivist. Feldman does not use the language of "standards," but rather "starting points" or "fundamental principles:"

A difficult project, which I will not undertake here, is to identify just what these starting points or fundamental principles might be and to explain how they might affect the sorts of disagreements under discussion. But whatever they are, I do not think that they will help solve the problem. Once people have engaged in a full discussion of the issues, their different starting points will be apparent. And then those claims will themselves be open for discussion and evaluation. ...Once you see that there are these alternative starting points, you need a reason to prefer one over the other. (Feldman, 2007: 206)

In this paragraph, Feldman proposes a principle for determining the rational validity of different starting points. Let's call this "Feldman's Principle":

(FP) If S has starting point α , and if S knows that there is an alternative starting point β , and if S is behaving rationally, then S has a reason to prefer α to β .

Notice that Feldman's Principle only applies when S knows about an alternative starting point. Thus formulated, (FP) is not strong enough to secure rational uniqueness, for in some cases, the agent may not know an alternative starting point exists. To secure Feldman's intended conclusion, we can strengthen the principle thus:

(S-FP) If S has starting point α , and if there exists another starting point β , and if S is behaving rationally, then S has a reason to prefer α to β .

If (S-FP) were true, then the permissivist could not appeal to alternative standards as an explanation of how the same set of evidence can be rationally interpreted as evidence for contradictory hypothesis. Fortunately for the permissivist, both (FP) and (S-FP) are false. Since (S-FP) logically entails (FP), we can restrict our attention to the falsity of (FP).

Rational reasoning must start somewhere. An infinite chain of reasoning might be possible, but it does not reflect our actual intellectual life. Circular reasoning is certainly possible, but just as certainly irrational. Even if our reasoning more closely models coherentism, there must be some initial starting points among which coherence can be found. These considerations are normally advanced on behalf of some version of foundationalism concerning propositional justification; however, they also reveal something about epistemic standards. According to what standards does one accept the epistemic standards one uses in weighing evidence? There must either be further standards by which those standards are evaluated or else those standards are in some way basic.

Suppose (FP) is true. If (FP) is true, then a rational agent must have a reason to prefer her standards to any other standard she knows to exist. This is a tall order indeed! For as we have seen, there is reason to think that there are an infinite number of potential standards. There are an infinite number of ways one could weigh the relative theoretic virtues of simplicity, explanatory power, and the like. Even if an agent is not cognizant of every alternative standard, the reflective person should quickly recognize how myriad and how disparate such standards might be.

But that something is hard is rarely a decisive objection against it. Perhaps rationality is simply very difficult. A deeper worry is that (FP) and uniqueness seem to commit one to broad skepticism. Suppose Kat and Michelle disagree about p. Kat believes that p and appeals to standards α in the evaluation of her evidence. Michelle considers the same evidence but believes that not-p according to her standards β . Applying (FP), each searches for a reason to prefer her own standards. Kat finds a reason through the application of meta-standard γ , whereas Michelle does the same via δ . They continue to disagree, and each becomes aware that yet another alternative starting point exists. They continue for some time. In fact, Kat and Michelle continue to exhaust all of their standards until they arrive at standards ψ and ω respectively. Neither can think of any additional reason or standard, unless it is according to itself (if this is

possible), by which they prefer ψ or ω ; they are basic, foundational standards. According to (FP), Kat and Michelle both ought to forfeit their standards and, accordingly, their beliefs about p .

This wouldn't be very problematic if it only meant forfeiting beliefs about p . Indeed, this is just what conciliationists, as well as many defenders of uniqueness, will recommend. What is problematic is that once one sees that there are alternatives even to her most basic and foundational standards, the threat of skepticism that began with forfeiting belief that p spreads to all other beliefs that are evaluated by the same standards. Suppose Kat and Michelle discover that they share a belief that q . Each of them evaluates q according to the same standards as she evaluated p . But now neither of them can continue to believe q —even though they agree about q —because it is licensed by the same standards they were already forced to abandon as a result of (FP).

Epistemic standards ought to remain consistent across domains. If Kat applied standards α in deliberation about p but some incompatible standards β in deliberation about q , we should think she is exhibiting a rational deficiency. Her standards are inconsistent. Of course, various features of some set of standards may be more salient when considering p than when considering q , but it would be surprising if the standards themselves substantially changed. Accordingly, what began looking like a relatively benign skepticism about p has revealed itself to be an infectious and widespread skepticism. At least, it will spread to all other beliefs that S arrived at by using the same standards as she used to arrive at the belief that p , regardless of whether there is any actual disagreement concerning the relevant belief. We may call this the problem of escalating skepticism.

Uniqueness may be preserved if the story about Kat and Michelle is impossible. If the ultimate, basic standards are such that they admit of no alternatives, then Kat and Michelle could never rationally arrive at the opposing, basic standards ψ and ω . In this case (FP) need not lead to a radically skeptical conclusion. I take it, however, that the onus is on the defender of (RUE) to show what these elusive standards might be. The history of human thought has not yet converged upon any such standards and appears to be in no danger of doing so.

4.3 Chancy Pill-popping and Epistemic Luck

White presents his own series of arguments intended to block the permissivist's appeal to alternative standards. White develops several stories involving someone who knows himself to be in an epistemically permissive situation. He is offered a pill that would change his standards so that he would go from rationally believing that p to rationally believing that not- p . The opportunity to take a standards-altering pill invites us to consider the apparent arbitrariness of having one set of standards as opposed to another—arbitrariness that seems to undermine the rationality of such standards in disputed cases. White writes:

I think we ought to be suspicious of this position [that different standards make permissivism plausible]... 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. (White, 2005: 451-452)

White assumes that the permissivist is committed to the following conditional: if one is in a situation that is known to be rationally permissive, one couldn't have arrived at one's standards in virtue of being rational. This is, according to White, because "it is consistent with [his] being rational that [he] adhere to rather different standards" according to which he would have believed differently. But it is unclear why the permissivist should accept this. Considering an analogy may be helpful. On at least some interpretations of quantum mechanics, the physical properties of subatomic particles do not determine their behavior; however, this is not to say that the activity of subatomic particles is not due to or in virtue of their physical properties. Similarly, someone could come to accept a set of standards in virtue of her rationality even if her rationality does not fully determine what standards she accepts.

Nevertheless, White is right that there does seem to be an important sense in which our pill-popping protagonist is extraordinarily lucky (if he gets the right result). Arriving at the correct standards appears to be a matter of luck. White thus raises an important question for the permissivist: If he is lucky, how can he be credited for a rational achievement? Simply showing that a belief is lucky, however, is not sufficient to demonstrate that it is irrational. One must show that the luck is malicious with respect to rationality and not benign.

Recent reflections on knowledge shed light on the distinction between malicious and benign luck. An underlying problem in nearly all Gettier cases is that the would-be knower seems to have benefited unduly from luck. Suppose I think I see a sheep in the field and justifiably form the belief, "There is a sheep in the field." Curiously, the sheep is actually a dog that just happens to look very much like a sheep, but the belief is true in virtue of a real sheep standing in the same field on the other side of a hill. My belief is true and justified, but the fact that I am right seems like dumb luck. I was lucky, epistemically speaking, that a real sheep happened to be somewhere else in the field. Our reactions to such cases reveal an anti-luck intuition about knowledge.

But not all luck is incompatible with knowledge. In Dahl's *Charlie and the Chocolate Factory*, Charlie finds one of five golden tickets distributed throughout the world that allows him to visit Mr. Wonka's marvelous chocolate factory. While there, he learns many strange things, among them this proposition: "There exists a chocolate waterfall." Charlie might think to himself, "Wow! I am so *lucky* to be here! If I hadn't happened to pick the fifth golden ticket, I would not be here now. I wouldn't have ever known that there was such a thing as a chocolate waterfall. In fact, I would have even had the false belief that there

weren't any chocolate waterfalls." Charlie's belief is indeed lucky in a very real sense. He could easily have had a false belief about the existence of chocolate waterfalls, and he would have failed to know if the world had been only slightly different. But the near-universal appraisal of such situations is that the relevant luck is epistemically benign. Someone can be lucky to find herself in a position to know without contradiction.³³

Here is another case that bears more directly on standards. Suppose that in a philosophical apocalypse, skeptics of the worst kind infiltrate the education system of every major government. Children are raised with epistemic standards that assign no evidential weight to information from the senses. You, however, grow up in post-apocalyptic Oxford, one of the few resistance holdouts against the new regime. Accordingly, you do assign evidential weight to the evidence of your senses. One day you think, "Wow! I am so *lucky* to have been born in Oxford! It is just a matter of chance that I was raised with the epistemic standards that allow me to admit the existence of an external world." Indeed, it is a matter of luck that you have the standards that allow you to have knowledge of the external world, but it is epistemically benign luck—and this despite the fact that you are even luckier in this case than White's pill-popping case in which the odds are presumably 50/50.

What is the difference between the luck in the Gettier case and Charlie's luck, or the luck exhibited by the denizens in post-apocalyptic Oxford? What explains why luck is malignant in the first case but benign in the others? This question deserves a paper in its own right, but for now let me gesture toward one answer.

Everyone was lucky to arrive at a true belief; however, only in the Gettier case was it lucky *given* the belief-forming method that the belief turned out to be true. Charlie was lucky to be in a position to form the belief, and the Oxonians were lucky to have an appropriate belief-forming method. But holding fixed their beliefs and belief-forming methods, there is nothing lucky about their having arrived at the correct answer. In the Gettier case, one is lucky to have arrived at a *true* belief; the luck comes "in between" the belief and the truth. In the benign cases, one is lucky to have arrived at a true *belief*; the luck does not come between the belief and the truth.

In response to White's charge, permissivists should maintain that any luck involved in acquiring truth-conducive standards is benign to rational belief. Rational belief involves correct calculation of the evidence given appropriate standards. If White's protagonist does not take the pill and correctly calculates the evidence upon what are, *ex hypothesi*, appropriate standards, his belief is rational. If those standards are truth-conducive, he may be lucky indeed, but not lucky in such a way that invalidates his rationality. He is lucky to have arrived at standards that produce true and rational beliefs; but, given the rational standards he in fact has, there is nothing lucky about those standards producing within him true and rational beliefs. The luck does not come between the belief and its rationality. Disagreement among evidential peers may indeed demonstrate something lucky (or unlucky) about our epistemic situation, but it need not demonstrate anything irrational.

³³I do not mean to imply that Charlie's beliefs are not "safe" in the technical, epistemic sense. The point is rather that certain kinds of luck (e.g. the luck of being in a position to know) are not relevant to the safety condition.

4.4 Remembering Standards: Assessing White's Arguments from Necessity and Reflection

Since his influential (2005) paper, White has published another article (2013) defending rational uniqueness—and it is probably the best defense of that thesis to date. Some of his arguments trade on the same concerns over luck and arbitrariness that we have already considered. Others, however, challenge permissivism in new and clever ways.

One especially persuasive argument draws attention to the Permissivist's violation of a very intuitive reflection principle.

Reflection If I know that tomorrow I will come to rationally believe P on the basis of new evidence without having lost any of my old evidence, then it is rational for me to believe P now. (White, 2013: 317)

The more one thinks about Reflection, the harder it is to deny. If the permissivist is to maintain her permissivism, she had better either be able to accept Reflection or else to show why denying Reflection is not as bad as it initially seems.

Are permissivists really committed to the denial of Reflection? White makes a compelling case that they are:

Instead of a P-pill and a not-P-pill we have a Truth pill and a Falsity pill. The Truth pill gives me a true belief regarding P and the Falsity pill gives me a false belief. (In either case the pill I take will cause me to base my belief on the available evidence.) You will flip a fair coin and feed me a pill depending on how the coin lands: Truth if the coin lands *heads*; false if it lands *tails*. (You will not let me know how the coin landed or which pill you have dropped in my mouth.) Suppose I come to believe P. Assuming that this is a permissive case my belief will be rational, and since I will know that P is the conclusion I have reached, I will also rationally believe that I have a true belief concerning P. But the same goes if I take the other pill and believe [not-P]: from my rational belief in [not-P] and my knowledge that I believe it I will rationally conclude that I have a true belief concerning P. So I can know in advance that regardless of which pill I take and whether P is true, upon examining the evidence and taking a pill I will rationally believe *now* that my future belief will be true. But I know it will be true only if the coin lands *heads*. It follows from strong permissivism and reflection that I can rationally predict in advance that this fair coin will land heads! But that is absurd. (White, 2013: 317)

Indeed, I believe that White is right and that the permissivist is forced to abandon Reflection. Despite the initial plausibility of Reflection, however, the permissivist can offer some principled reasons to reject Reflection in favor of a more nuanced principle.

Recall that the permissivist—at least the permissivist as depicted in this paper—believes that rationality is (at least) a three-place relation. Rationality is not just a function of evidence to doxastic attitudes. The rationality of a doxastic attitude toward a hypothesis is a function of the evidence *and* the standards according to which the evidence is evaluated for that hypothesis.

The permissivist notes that White’s Reflection principle includes a clause that accounts for the possibility of lost evidence; the permissivist would also like to include a clause that accounts for the possibility of changed standards:

Standards-Sensitive Reflection If I know that tomorrow I will come to rationally believe P on the basis of new evidence without having lost any of my old evidence *and without having changed epistemic standards*, then it is rational for me to believe P now.

It matters very much *how* the truth and falsity pills in White’s story operate. Suppose they operate without altering the agent’s standards. If such a story were possible, the permissivist would be forced to abandon even Standards-Sensitive Reflection. In my view, the permissivist should agree with White that this really is absurd. White’s argument has shown us something very interesting about the permissiveness of evidence. Evidence cannot be permissive (in the strong sense)³⁴ absent different standards.

Suppose, however, that the pills in White’s story operate by altering the standards of the agents. Then all bets are off. The change in standards disrupts our ability to make stable links between our future and current beliefs. Yes, the permissivist must reject Reflection; she needn’t be concerned, however, since she can maintain Standards-Sensitive Reflection.

A second, new argument from White’s recent article concerns the allegedly necessary connection between total evidence and the hypotheses that evidence supports: “if E supports p then necessarily E supports p” (White, 2013: 314). When conjoined with the premises that evidence is unidirectional (see 4.1) and that rationality requires one’s conclusions to be supported by the evidence, it constitutes a valid argument against permissivism. White presents two considerations in support of this claim. The first concerns our ability to know what our evidence supports. If a total set of evidence supported a conclusion contingently, “it would be unclear how we can access what our evidence supports” (White, 2013: 314). After all, how would we *know* that the evidence supported the conclusion it *actually* supports if it *could have* supported something else instead? White’s second objection is that such contingency makes it possible to tamper with what is rational without tampering with the evidence. White tells the following story:

Our total evidence E does in fact support the standard thesis C of climate change. But now suppose that that very evidence need not have supported this conclusion. If this were so then it would seem that industry lobbyists could in principle manipulate the world such that E supports not-C instead. Hence they could make it the case that we can rationally *believe* that carbon emissions do no harm (and rationally act on this!) without having done anything to prevent disastrous climate change. This seems absurd. (White, 2013: 314).

Regarding White’s first concern about how the permissivist would know which hypotheses the evidence supports for her since (according to the permissivist) such support relations are contingent, the permissivist has a simple response: we can know what our evidence *actually* supports by considering our evidence

³⁴The argument from efficient evidence still shows that when interests differ, evidential peers may arrive at different—although not contrary—beliefs.

in light of the standards we *actually* have.³⁵ On the account this paper has sketched, the only thing that the evidential support relations are dependent upon is the agent's standards. In many cases, an agent can simply reflect on her own standards in order to have a sense of what they are. Once the permissivist has controlled for her standards, the permissivist should be able to assess the support relations (for her) between sets of evidence and hypotheses as well as anyone. After all, the evidential support relations are not dependent on anything else.

Admittedly, there may be cases in which an agent's own standards are not perspicuous to her. Perhaps a better response is to point out that the metaphysical necessity of evidential support relations wouldn't obviously do any work toward explaining our epistemic ability to assess what is rational. It is metaphysically necessary that water=H₂O. But it does not seem that it would have been any harder (or easier) for scientists to discover the formula for water if, *per impossible*, the relation had been contingent. True, if we knew that evidence necessarily supported its conclusions, we would never need to ask, "why believe that the evidence supports C for me though it might (metaphysically) support not-C for me?" But it would not keep us from asking, "why believe that the evidence supports C though it might (epistemically, as far as I know) support not-C?" And this seems like just as difficult a question.

Concerning White's story about the industry lobbyists, it once again, as in the pills case, depends a great deal on *how* the industry lobbyists are accomplishing their task. Are they doing it by—somehow—altering the internal standards by which we assess the evidence? If so, the permissivist should accept that the lobbyists really can change what is rational for us to believe without lifting a finger to prevent climate change or otherwise to alter our evidence. But once we imagine that the lobbyists have the power to alter our internal standards, it ceases to seem absurd that they can change what is rational for us to believe. We have given them powers akin to those of Descartes's evil demon! It is an unfortunate feature of our epistemic life that we can be so fooled, but it is also an unavoidable one.

On the other hand, if the lobbyists are supposed to change what is rational to believe without either changing the evidence or the internal standards by which the evidence is judged, the permissivist may readily agree with White that such would be truly absurd. That is not the way that evidence supports its conclusions contingently. It only supports its conclusions contingently insofar as it depends on the actual standards of the agent who is forming the judgment.

4.5 Conclusion

White and Feldman force the permissivist to draw careful distinctions. When speaking precisely, we must always remember that evidence is a function not only of evidence but also of an agent's epistemic standards—standards that include, at least, the way an agent weighs the theoretic virtues against each other. When confronted with the luck involved in permissive circumstances, the permissivist must be prepared to judge between benign and malicious cases. Yet

³⁵I do not mean to imply that in order to know whether one's evidence supports a given hypothesis one must *know* what standards one has, only that one can (at least sometimes) know what one's evidence supports by evaluating the evidence from a perspective that includes certain standards.

even with these distinctions in mind, questions remain. Does it matter whether one knows (or rationally believes) that one is in an epistemically permissive situation? Would knowledge of such second-order evidence defeat the first-order beliefs licensed by permissive standards or, at least sometimes, leave them intact? Is epistemic permissiveness robust enough to withstand sustained peer disagreement?

Sadly, these questions will not find even an attempted answer here. In this respect, the goals of this paper have been somewhat modest. The first goal has been to distinguish between several kinds of permissivisms in order to clarify the terminology of an ongoing debate. The second has been to present and defend a positive case for extreme permissivism. In this endeavor, I am indebted especially to Nelson and Douven, whose writings have permeated my thinking on this subject. The argument from efficient evidence sought to show that agents are not always required to endorse hypotheses that are evidentially supported for them. The argument from abduction contended that fully rational agents may sometimes endorse contradictory hypotheses when they do not weigh their evidence according to the theoretic virtues in the same way. Rational uniqueness could be preserved if there were a uniquely best way to weigh the theoretic virtues, but no such optimum can be found. The final goal has been to scrutinize the most persuasive arguments offered for uniqueness and to offer new responses on behalf of the permissivist. White's and Feldman's arguments deserve serious consideration; however, they do not compel the permissivist to abandon her position.

Evidential peers may sometimes arrive at very different conclusions without having made any mistakes in reasoning. Even the greatest minds do not always think alike.

Acknowledgments

I owe special thanks to Max Baker-Hyitch, Laura Callahan, Joshua Layton-Wood, Brian Leftow, Tim Mawson, Andrew Moon, Mark Nelson, Danielle Willard-Kyle and three anonymous reviewers for their helpful comments on drafts of this paper. Each of them is truly a great mind.

References

- [1] Alston, William P. 1980. "Level Confusions in Epistemology." *Midwest Studies in Philosophy* 5: 135–50.
- [2] Baker-Hyitch, Max. 2014. "Religious Diversity and Epistemic Luck." *International Journal for Philosophy of Religion*. 76 (2): 171-191.
- [3] Ballantyne, Nathan, and E.J. Coffman. 2011. "Uniqueness, Evidence, and Rationality." *Philosophers' Imprint* 11 (18).
- [4] ———. 2012. "Conciliationism and Uniqueness." *Australasian Journal of Philosophy* 90: 657–70.
- [5] Bonjour, Laurence. 2002. *Epistemology: Classic Problems and Contemporary Responses*. Rowman & Littlefield.

- [6] Chisholm, R. M. 1977. *Theory of Knowledge*. 2nd ed. Upper Saddle River, NJ: Prentice-Hall.
- [7] Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2): 187–217.
- [8] ———. 2013. "Epistemic Modesty Defended." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- [9] Christensen, David, and Lackey, eds. 2013. *The Epistemology of Disagreement: New Essays*. Oxford University Press.
- [10] Clifford, William K. 1886. *Lectures and Essays*. Edited by Leslie Stephen and Frederick Pollock. London: Macmillan and Co.
- [11] Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Clarendon Press.
- [12] Dougherty, Trent, ed. 2011. *Evidentialism and Its Discontents*. Oxford University Press.
- [13] Douven, Igor. 2009. "Uniqueness Revisited." *American Philosophical Quarterly* 46: 347–61.
- [14] ———. 2011. "Abduction." Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/spr2011/entries/abduction>.
- [15] Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60: 667–95.
- [16] ———. 2004. "Having Evidence." In *Evidentialism: Essays in Epistemology*, edited by Earl Conee and Richard Feldman, Oxford University Press.
- [17] ———. 2006. "Epistemological Puzzles About Disagreement." In *Epistemology Futures*, edited by Stephen Heather. Oxford University Press.
- [18] ———. 2007. "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, edited by Louise Antony. Oxford University Press.
- [19] ———. 2009. "Evidentialism, Higher-Order Evidence, and Disagreement." *Episteme* 6 (3): 294–312.
- [20] Feldman, Richard, and Ted A. Warfield. 2010. *Disagreement*. Oxford University Press.
- [21] Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121–23.
- [22] Hawthorne, John, and Amia Srinivasan. 2013. "Disagreement Without Transparency: Some Bleak Thoughts." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.

- [23] James, William. 1912. *The Will to Believe and Other Essays in Popular Philosophy*. Longmans, Green, and Co.
- [24] Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology*, edited by John Hawthorne and Tamar Szabó Gendler, 1:167–96. Oxford University Press.
- [25] ———. 2006. "Evidence." Edited by Edward N. Zalta. *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/evidence/>.
- [26] ———. 2007. "Peer Disagreement and Higher Order Evidence." In *Disagreement*, edited by Richard Feldman and Ted Warfield, 2010:111–74.
- [27] ———. 2008. "Disagreement, Dogmatism, and Belief Polarization." *The Journal of Philosophy* 105 (10): 611–33.
- [28] ———. 2011. "Following the Argument Where It Leads." *Philosophical Studies* 154 (1): 105–24.
- [29] ———. 2013. "Evidence Can Be Permissive." In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, 2nd ed. Blackwell.
- [30] Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1: 337–55.
- [31] Kuhn, Thomas. 2012. *The Structure of Scientific Revolutions*. 3rd ed. Chicago University Press.
- [32] Kvanvig, Jonathan L. 2013. "Perspectivalism and Reflective Ascent." In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford University Press.
- [33] ———. 2014. *Rationality and Reflection: How to Think about What to Think*. Oxford University Press.
- [34] Lee, Matthew. 2013. "Conciliationism without Uniqueness." *Grazer Philosophische Studien* 88.
- [35] MacFarlane, John. 2007. "Relativism and Disagreement." *Philosophical Studies* 132: 17–31.
- [36] Nelson, Mark T. 2005. "Telling It like It Is: Philosophy as Descriptive Manifestation." *American Philosophical Quarterly* 42: 141–56.
- [37] ———. 2010. "We Have No Positive Epistemic Duties." *Mind* 119: 83–102.
- [38] Pittard, John. 2014. "Conciliationism and Religious Disagreement." In *Challenges to Moral and Religious Belief: Disagreement and Evolution*, edited by Michael Bergmann and Patrick Kain. Oxford University Press.
- [39] Rosen, Gideon. 2001. "Nominalism, Naturalism, Epistemic Relativism." *Philosophical Perspectives* 15: 60–91.
- [40] Stapleford, Scott. 2013. "Imperfect Epistemic Duties and the Justificational Fecundity of Evidence." *Synthese* 190(18):4065–4075.

- [41] Van Inwagen, Peter. 1996. "It Is Wrong, Everywhere, Always, and for Anyone, to Believe Anything on Insufficient Evidence." In *Faith, Freedom, and Rationality: Philosophy of Religion Today*, edited by Jeff Jordan and Daniel Howard-Snyder, 137–53. Rowman & Littlefield.
- [42] ———. 2009. "Listening to Clifford's Ghost." *Royal Institute of Philosophy Supplement* 84: 15–35.
- [43] Van Wietmarschen, Han. 2013. "Peer Disagreement, Evidence, and Well-Foundedness." *The Philosophical Review* 122 (3): 395–425.
- [44] Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford University Press.
- [45] White, Roger. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19: 445–49.
- [46] ———. 2009. "On Treating Oneself and Others as Thermometers." *Episteme* 6 (3): 233–50.
- [47] ———. 2013. "Evidence Cannot Be Permissive." In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, 2nd ed. Blackwell.
- [48] Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.