episteme

ARTICLE

A New Argument for Uniqueness about Evidential Support

Paul Forrester

Yale University, New Haven, Connecticut 06520, USA
Email: paul.forrester@yale.edu

Abstract

In this paper I identify a family of explanatory demands facing permissivists, those who deny the uniqueness thesis, according to which every body of evidence rationally permits exactly one doxastic attitude for a person to have in light of that evidence. Call a pair of a body of evidence and a proposition a *permissive case* just in case there is more than one attitude that is permitted for someone who has that body of evidence to take to that proposition. Uniquers claim that there are no permissive cases, and permissivists deny this. The uniqueness thesis is a strong claim, and is vulnerable to counterexamples because it is a universal generalization, and indeed, many permissivists argue for their view by identifying putative counterexamples to the uniqueness thesis. However, in virtue of advancing these putative counterexamples, permissivists incur explanatory demands. If some, but not all, cases are permissive cases, then permissivists owe us an explanation of why some cases are permissive and others are not. Likewise, for each permissive case, permissivists must explain why some attitudes towards the proposition are permitted and others are not. Permissivists draw arbitrary lines between permissive and impermissive cases, and between permitted and impermitted attitudes, giving rise to distinctions which need explaining.

I shall argue that permissivists cannot discharge these explanatory burdens in a satisfying way. After carefully presenting these explanatory demands in section 2, I consider how permissivists might answer them in section 3. I argue that the only permissivists who are able to successfully answer the explanatory demands are extreme permissivists like subjective Bayesians. Most philosophers, however, will find this epistemological outlook implausible because it contains no substantive constraints on rational belief. I also show that previous attempts by permissivists to explain why there are permissive cases merely relocate, rather than answer, the explanatory demands I have identified. The final section argues that permissivists cannot soften up the explanatory demands, either. I show that permissivist cannot appeal to the vagueness of the notion of a permissive case, or of an attitude's being permitted, in order to make the line between the permitted and the impermitted attitudes appear less arbitrary. Nor, I argue, can the permissivist hold that our ignorance of which cases are permissive, and how permissive they are, lessens the explanatory burdens they face.

Unfortunately, permissivists have been most interested in attempting to identify a counterexample to the uniqueness thesis, and have not been concerned with the explanatory demands which they incur in so doing. If my arguments are successful, then, permissivism is a much less plausible view than it is currently given credit for, and, at the very least, permissivists have a lot of work to do in further articulating and generalizing their view so as to provide an adequate explanation of the scope of epistemic permission.

1. Introduction

The debate between uniquers and permissivists concerns whether a body of evidence must rationalize just one, or could rationalize more than one, doxastic attitude to a given proposition, respectively.[1] A proposition–evidence pair for which there are two or more doxastic attitudes which are permitted for a person to have in light of that evidence is called a *permissive case*. Permissivism is the view that there is at least one permissive case, and the uniqueness thesis is the claim that there are no permissive cases.

Since the uniqueness thesis is a universal generalization, a common strategy among permissivists is to identify a single counterexample to it; such a counterexample would falsify the thesis and show that permissivism was true. Permissivists allege that there are a variety of situations in which two persons with the same evidence rationally form different doxastic attitudes on the basis of that evidence. Rosen (2001), for example, holds that two members of a jury who are presented with the same evidence in a trial may nevertheless come to differing opinions concerning the guilt of the defendant, though neither one is irrational. Other examples abound.[2] Regardless of the merits of any particular putative counterexample, permissivists have pointed out that the uniqueness thesis is a very strong claim since it is a universal generalization. It takes just one case where rationality is permissive for permissivism to be true, but for uniqueness to be true, *all* cases must *not* be permissive. As such, it would take incredibly strong arguments for uniqueness and a very plausible error theory about the intuitions motivating the putative counterexamples to convince us that there was not a single permissive case.[3] Permissivism, it is alleged, should be the default view.

Uniquers have several strategies for responding to such putative counterexamples. First, they could hold that the parties in question do not in fact have the same evidence: there is at least one piece of evidence that one party has that the other lacks which rationalizes their difference in opinions.[4] Second, uniquers could hold that there is in fact only one attitude rationalized by the evidence, but it is an imprecise one which includes both more specific attitudes which are had by the different persons.[5] Finally, uniquers could hold that one party has simply made a mistake in reasoning and evaluating the evidence and thus has a less than fully rational attitude.[6]

---

[1]White (2005).

[2]Schoenfield (2014) suggests that belief that God exists forms a permissive case, on at least some bodies of evidence. Titelbaum and Kopec's (ms) reasoning room is another putative counterexample to uniqueness.

[3]This sentiment is on full display in Kelly (2013).

[4]Goldman (2010).

[5]See Castro and Hart (2019) for a critical evaluation of this strategy.

[6]There are also meta-epistemological arguments for uniqueness, on which see Greco and Hedden (2016).

However, uniquers have a more general and theoretically well-motivated reply to the threat of counterexample, one which it is this paper's central purpose to develop. Permissivists are correct to assert that uniqueness is a strong view because it is a universally quantified claim. But permissivists advance a *merely existentially quantified claim*, according to which some, but not all, cases are permissive cases. Though philosophers advancing universally quantified claims are vulnerable to counterexample, philosophers advancing merely existentially quantified claims incur an explanatory burden. They must explain why the predicate in question is true of some things but false of others – in virtue of what are some things F and others not-F? In the case at hand, permissivists need to explain the scope of epistemic permission – they have to explain why some cases are permissive and others are not. I shall argue that permissivists cannot discharge this burden in a satisfying way.

My arguments in this regard fit quite nicely with a few other anti-permissivist lines of thought in the recent literature. In a permissive case, the permissivist is permitted to the existence of a range of credences which are permissible, though credences outside that range are not permissible. Schultheis (2018) argues that since we do not know precisely where the edges of the permissive range lie, there could not be such a range. Horowitz (2014) argues that the permissivist's commitment to a range of permissible credences is inconsistent with an accuracy-based explanation of the value of rationality. I add a third argument against the existence of a permissible range: the permissivist is required to give an explanation of the extent of the permissible range of credences, but they cannot do so. I will discuss these arguments and their relationship with my arguments at greater length below. Together, these three arguments constitute a decisive case against the idea that there is sometimes a permissible range of attitudes that one can have in light of one's evidence.

Here is how we will proceed. In section 2, I will more precisely articulate the explanatory demands which face the permissivist. In the second, I will consider how they might be answered, and ultimately conclude that the only permissive views that avoid these explanatory demands are very implausible. Finally, I will consider how the explanatory demands might be softened, by appealing to vagueness or our ignorance of the facts about which cases are permissive, concluding that neither strategy ultimately pans out. Hence, if my arguments are successful, permissivism is a much less plausible view than it is currently given credit for, and, at the very least, permissivists have a lot of work to do in further articulating and generalizing their view so as to provide an adequate explanation of the scope of epistemic permissions.

2. Explanatory demands facing the permissivist

A few terminological notes are in order before we can state the explanatory demands facing permissivism. I will distinguish four kinds of cases, where a case is a triple of a person, S, their body of evidence, E, and a proposition, P. An *impermissive case* is one in which only one doxastic attitude is permitted for S to have to P. An *epistemic dilemma* is a case where there are no attitudes that S is permitted to have to P. A *permissive case* is one in which there is more than one doxastic attitude that S is permitted to have to P. Permissive cases come in two kinds: wide and narrow.

A *wide permissive case* is one where the individual is permitted to have any doxastic attitude to P, and a *narrow permissive case* is a permissive case that is not wide.[7]

Formally, we can think of the uniqueness thesis this way. Consider the following three sets:

A: the class of all doxastic attitudes.
P: the class of all propositions.
E: the class of all bodies of evidence.

The uniqueness thesis is the claim that epistemic rationality can be formally represented as a function from $E \times P$ (the Cartesian product of E and P) into A. Each pair of a body of evidence and a proposition corresponds to exactly one doxastic attitude – the one that the person is permitted to take to the proposition given her evidence. This function encodes all of the information about which states are epistemically permitted in response to what evidence. Permissivism is the thesis that epistemic rationality is not a function, but merely a mapping from $E \times P$ to A, for a pair of E_i and P_i could be mapped into more than one element of A (this is a permissive case).

One more clarification is in order in light of the above characterization of the terrain. Many epistemologists hold that one can have both beliefs and credences towards propositions.[8] For example, one might rationally have a belief that one will lose the lottery and rationally have a credence of 0.99999 that one will lose the lottery. Anyone who holds this view will count as a permissivist, in the sense that I have identified above, as more than one doxastic attitude is permitted by the person's evidence (if we call both credences and belief-like attitudes "doxastic attitudes").

These views should not be trivially classified as versions of permissivism, though. We should instead follow Jackson (2021) in distinguishing credal uniqueness and belief uniqueness. According to the latter, a person's evidence supports exactly one of the belief-like attitudes: belief, suspension and disbelief. And according to the former, a person's evidence supports exactly one of the credal attitudes: credences in [0, 1]. Of course, one of these uniqueness theses may be true and the other false.[9] More generally, we could say that, for each attitude type, a person's body of evidence rationalizes exactly one attitude of that type. This characterization is entirely neutral with regard to views about the nature of the different types of doxastic attitudes, with the caveat, which I shall assume throughout the paper, that for each attitude type, a person can only have one attitude of that type to a proposition at a time (e.g. someone could not both believe P at t and suspend judgment on P at t, even if, according to permissivism, both such attitudes are permitted). Philosophers who believe that there is more than one kind of doxastic attitude should be satisfied with this reformulation of the

---

[7]Some authors, e.g., Feldman (2006) and Matheson (2011), formulate the uniqueness thesis as the claim that in every case, at most one attitude is permissible for one to have in light of one's evidence (rather than "exactly one"). This formulation groups impermissive cases with epistemic dilemmas, and it is unclear why these should go together in the formulation of the uniqueness thesis. Indeed, I shall argue below that epistemic dilemmas are permissive cases (given suitable assumptions) so this gives us reason to use the "exactly one" formulation of the view rather than the "at most one" formulation.

[8]For a survey of views about the relationship between beliefs and credences, see Jackson (2021).

[9]It is often pointed out (e.g. in Kelly 2013) that the uniqueness thesis is much more plausible when it is taken as a thesis about coarse-grained attitudes rather than fine-grained ones, because it is more difficult to see how each body of evidence could support exactly one of an infinite number of credences, rather than exactly one of three belief-like states.

uniqueness thesis, as it does not trivially classify them as permissivists, and it allows uniqueness to be true of some attitude types and false of others. I think that this distinction allows us to conduct the debate about uniqueness in a way that entirely cross cuts debates about the various kinds of doxastic attitudes.

My arguments for uniqueness are quite general, and will support the version of the uniqueness thesis for any attitude type. This is because they do not concern features particular to any attitude type. Rather, they concern explanatory demands facing anyone who posits a distinction between permissive and impermissive cases. Let us turn to those arguments.

One who advances a universally quantified claim incurs a substantial commitment: that everything (within the scope of the quantifier) is a certain way. Claiming that everything is F commits one to a thesis that is easy to disprove (just find one thing that is not F!) and that is difficult to justify, for in order to establish the thesis, one must examine everything that there is, showing that each one is F. One must defend against counterexample and pass an extremely high bar of justification if one is advancing a universally quantified claim.

By contrast, a *merely existentially quantified claim* – a claim of the form: some, but not all, things are F – seems not to require such a high bar of justification. All one must do to establish the truth of the claim is find one thing that is F, and another thing that is not F. One need not examine everything within the scope of the quantifier, like one advancing a universally quantified claim. Nor is one vulnerable to counterexample, since once examples of something that is F and something that is not F are found, no further examples could falsify the claim.

Uniquers advance a universally quantified claim: for every pair of a body of evidence and a proposition, a rational agent with that body of evidence is permitted to have only one doxastic attitude to that proposition. In other words, there are no permissive cases – every case is an impermissive case. Permissivists advance an existentially quantified claim, that there are permissive cases: sometimes, a person's body of evidence rationalizes more than one doxastic attitude to have to a given proposition. Most permissivists will advance a merely existentially quantified claim – for they will not think that every case is a permissive one. We will consider the *extreme* form of permissivism on which every case is permissive in section 3. This section will mainly concern those *moderate* permissivists who advance a merely existentially quantified claim about epistemic rationality.

Some permissivists have thought that these considerations constitute an argument against uniqueness: uniquers, they say, must defend a very strong claim since it is universally quantified, but permissivists defend a very weak claim. There may only be a few permissive cases, but even that would be sufficient to show that permissivism was true. Right off the bat, then, permissivism is a much more plausible and defensible thesis than uniqueness.[10]

This seems right, and indeed, from what I have said so far, it might be judged that it is unwise to ever advance universally quantified claims. But we can now identify an equally weighty downside of advancing a merely existentially quantified claim. The person advancing a merely existentially quantified claim incurs an explanatory burden: if something is F and something else is not F, then why are some things F and other things not F? In virtue of what are some things F and others not? What criterion could we use to sort the things into those that are F and those that are not? The person advancing the

---

[10]Kelly (2013).

merely existentially quantified claim advances a logically weaker thesis, but in so doing they incur a heavy explanatory burden. They must give us an explanation of the distinction they posit between the Fs and the non-Fs.

Sometimes, this burden is easily discharged. Some, but not all, cars are blue. I can establish this claim by identifying two examples. I can discharge the explanatory burden which I have incurred by giving a causal explanation of why some cars are blue and others are not: the manufacturers and owners of cars paint some blue, and paint others different colors. This explanation is perfectly satisfactory, and needn't call into question any more fundamental issues in the philosophy of color or concerning the vagueness of "blue" or "car." The explanatory burden here is so trivial to discharge that we would not even recognize it unless it is pointed out.

In more theoretical or fundamental discussions – and most philosophical topics fall in this category – however, the explanatory demands are not as easily discharged. Consider the following claim: some, but not all, persons have rights. Many philosophers will want to advance such a claim. Such philosophers may think that fetuses, the senile, or those who have committed heinous crimes, are persons yet do not have certain rights. Such a philosopher incurs an explanatory burden which is much more difficult to discharge than the one in the car example above. The philosopher must explain at exactly what point in the developmental process a person begins to have rights, and why this physical fact about the developmental process is normatively connected to the fact about what rights the person has. Or the philosopher must explain exactly how heinous a crime must be in order for the person who committed it to forfeit their rights, and exactly what process of restitution must be engaged in for the forfeited rights to be restored, if this is possible. In sum, the philosopher will have to develop a complicated theory to explain why some but not all persons have rights. What's more, for such a philosopher, the notion of being a person can do no explanatory work for showing why an entity has rights, since many persons do not have rights.

Other philosophers maintain that all people have rights. Any of the cases discussed in the last paragraph will trouble this philosopher, too, this time as putative counterexamples to the claim that all people have rights, rather than cases which make salient the explanatory burden which faces the holder of the merely existentially quantified claim. There are two strategies for dealing with these cases *qua* putative counterexamples. First, the philosopher could deny that the relevant entities are persons. Perhaps a human organism is not a person until it is born, or a human organism ceases to be a person upon the commission of some heinous crime, etc. These are options to the philosopher if she wants to deny that such entities have rights and still maintain her claim that all persons have rights. For this philosopher, the notion of being a person does work in the explanation of why something has rights. For this philosopher, the difficult questions about which entities have rights are then turned into (perhaps less difficult) questions about which entities are persons. Alternatively, the philosopher can hold that the human organisms in the examples are indeed persons, and as such, do have rights after all.

In fact, justifying a universally quantified claim may in some cases prove easier than answering the explanatory demands facing a merely existentially quantified claim. This is because the philosopher advancing a universally quantified claim may posit a conceptual or constitutive connection between personhood and being a rights-bearer, or whatever properties are in question. If the philosopher is in a position to deduce, *a priori*, from the notion of a person that they are rights-bearers, this can serve as a defense against counterexamples (remember, the philosopher still has the two strategies

mentioned above). Such an argument would be more than adequate to meet the high bar for the justification required by universally quantified claims. More generally, this conceptual or constitutive connection strategy is available to anyone who makes universally quantified claims that they take to be justified *a priori*, and there is nothing preventing uniquers from availing themselves of this strategy.

Let us now examine how these considerations bear on the debate over the uniqueness thesis. Permissivists must explain why some bodies of evidence are such that they rationalize more than one doxastic attitude toward a proposition, and others are not. For permissivists, all cases fall into two classes, permissive cases and impermissive cases, and the permissivist must explain why a given case falls into one class rather than the other. Permissivists must give us substantive necessary and sufficient conditions a case must satisfy in order to be permissive.

Uniquers, on the other hand, hold that every body of evidence is such that there is exactly one permitted attitude in each case. For uniquers, all cases belong to one category: impermissive cases. So, uniquers incur no explanatory burden. Extreme permissivists, who think that every case is a permissive case, also have no explaining here to do, for they think that all cases also belong to one category: permissive cases.

An additional explanatory burden besets permissivists, both moderate and extreme: they must explain why, in each permissive case, some doxastic attitudes are permitted in that case, and others are not permitted in that case. In wide permissive cases, this explanatory burden does not arise, since in such cases all doxastic attitudes are on a par – they are all permitted. In a narrow permissive case, by contrast, the permissivist sorts all of the doxastic attitudes into two categories: the ones that are permitted in the case, and the ones that are not permitted in the case. For example, a permissivist might want to say that given evidence E, a rational agent is permitted to have credences in [0.6, 0.8] in proposition P, and not permitted to have any other credence in P. The permissivist must explain why the permissible range is [0.6, 0.8] rather than, say, [0.6, 0.8 + ε]. What is the important difference between credences of 0.8 and 0.8 + ε that explains why the former credence is permitted by epistemic rationality and the latter credence is forbidden? The permissivist must explain the scope of epistemic permission not only across the cases, but also within each case. Otherwise, the sharp borders of the permissible range are going to be completely arbitrary.[11]

To put the point more sharply, permissivists criticize uniquers because permissivists think that some bodies of evidence do not contain enough information to rationalize exactly one out of a wide variety of doxastic attitudes to a proposition. Yet, permissivists are not in a dialectical position to advance this kind of criticism. This is because permissivists think that bodies of evidence contain enough information to determine two points at which there are sharp transitions from impermissible to permissible credences.[12] Uniquers are right to wonder how the permissivist thinks that the evidence could contain enough information to determine a precise range of permissible

---

[11]I should note that this problem is not just one for the permissivist about credences. Since there are three belief-like attitudes – belief, suspension, and disbelief – permissive cases are either ones in which only two of those attitudes are permitted (narrow), or ones in which all three are permitted (wide). So, permissivists must explain the range of permission in narrow permissive cases even for coarse-grained attitudes.

[12]I assume that there is just one continuous zone of permitted credences, because I assume that permissivists will want to say that any linear combination of permitted credences is a permitted credence. Intuitively, it would be odd if, say, [0.5, 0.6] and [0.8, 0.8001] were the credences permitted by E to P, and all others were not permitted. In any case, failure to hold that the permissive zone is continuous

credences, but somehow could not contain enough information to determine a single uniquely rational credence. All bodies of evidence have the ability to precisely determine which attitudes are permissible and which impermissible, even from the permissivist's perspective. This leaves us wondering why permissivists insist that more than one attitude is permitted in some cases because of the informational content of some bodies of evidence. Thus, once we see this point, the common permissivist line of thought that permissive cases are a result of the fact that we sometimes have sparse or imprecise evidence loses its force, since even permissivists must say that such bodies of evidence determine a sharp cutoff between impermissive and permissive responses. Permissivists might try to soften up these explanatory demands by saying that the edges of the permissive zone are metaphysically indeterminate, or that we simply are never in a position to know where they are. We will evaluate these responses in section 4.

There is a final explanatory burden: if the permissivist is committed to the existence of both wide and narrow permissive cases, she must explain why some permissive cases are wide, and others narrow. Such a permissivist acknowledges two distinct kinds of permissive cases, wide and narrow, and needs to give us an explanation of the difference between the narrow and wide permissive cases.

Thus, permissivists need a theory of epistemic rationality which would enable them to answer three questions about the scope of epistemic permissions:

1. Why are some cases permissive and some cases impermissive?
2. In narrow permissive cases, why are some doxastic attitudes permitted and others not?
3. Among the permissive cases, why are some wide and others narrow?

Such a theory would allow the permissivist to discharge the explanatory burdens they incur in virtue of advancing a merely existentially quantified claim. The only permissivist theory that dodges these explanatory demands entirely is what I call "nihilist permissivism." If nihilist permissivism is true, then for every body of evidence E, person S, proposition P and doxastic attitude D: when she possesses E, S is permitted to take D to P. In other words, no attitude is impermissible – every case is a wide permissive case. Since there is no divide between permissible and impermissible attitudes, there is no distinction that calls for explanation. I find this view very implausible, despite the fact that it evades the central argument of my paper.[13] Any non-nihilist form of permissivism posits some kind of structure to the space of epistemic requirements and permissions, a structure that requires a commensurate level of explanatory support. Another implausible view that I will reject is that these distinctions are brute: there is no explanation of the extent of epistemic permission. The difference between permissive and non-permissive cases is something that cries out for explanation, and to be told that there is no explanation is just as bad as to ignore the question entirely.

Unfortunately, most of the papers arguing for permissivism simply seek to provide a counterexample to the uniqueness thesis, and do not even recognize that they incur these explanatory burdens, much less try to answer them. The lesson of our analysis is that coming up with an intuitively plausible counterexample is not sufficient to

---

would simply multiply the explanatory burdens that the permissivist faces, committing them to more sharp transitions between permission and impermission.

[13]I am not aware of any epistemologist who holds this view, with the possible exception of Schoenfield (2022).

convince us that we should give up a universally quantified claim. This is not sufficient to shift the balance of considerations in favor of a merely existentially quantified claim against a universally quantified claim. It is an important step, no doubt, but the counterexample-provider also has a distinction to explain: *why* the predicate is true of some things, but false of others. In the absence of such an explanation, we may indeed be better advised to reject their counterexample and the intuitions behind it rather than accept an account with unexplained and *ad hoc* distinctions. Philosophers generally have not recognized that advancing a counterexample incurs an explanatory burden, and accordingly, permissivists in particular have not taken up the challenge of explaining when and why epistemic permission arises.[14] In the next section I will explore a couple of avenues on the permissivist's behalf, ultimately rejecting both.

I should note before moving on that the permissivist's commitment to the existence of a range of permitted attitudes is the source of problems other than the ones that I am pressing, some of which have been highlighted in the recent literature. Schultheis (2018) argues that we could not know exactly where the endpoints of the permissible range lie because of our limited ability to ascertain what the requirements of rationality are (compare: we can't tell what the height of a tree is down to the nearest inch just by looking, because our perceptual faculties are not perfect, and likewise, we couldn't tell exactly what credences are permitted in a proposition down to a 0.000001 interval). Schultheis then argues that if one cannot know where the edges of the permissible range are, then it is impermissible for one to have a credence that is in fact the endpoint of the permissible range. This is because if one is choosing between two options, one of which one knows to be permissible and another of which might be permissible or might not, then one should choose the option which one knows to be permissible. But this contradicts the assumption that the attitude on the edge of the permissible range is indeed permissible. Hence, there is no such permissible range.

Likewise, Horowitz (2014) argues that permissivism is incompatible with an accuracy-first explanation of the value of rationality. Consider a credence A near the upper edge of the permissible range of credences. From the perspective of A, a credence B that is just a bit higher than A but outside of the permissible range is more expectedly accurate than a credence C further away from B than A, but in the permissible range near the bottom. This means that for the permissivist, rationality is not tightly connected with expected accuracy. Permissivists, if they do not have credences in the dead center of the permissible range, will always regard some rationally permissible credences as less expectedly accurate than some impermissible ones. And in Horowitz's view, this is a damning objection to permissivism.

---

[14]This is not to say that counterexamples to generalizations are never correct. Indeed, they often are – my grey car refutes the generalization that all cars are blue. In a more theoretical domain, Gettier cases refute the K = JTB analysis of knowledge. But in order to be sure that the putative counterexamples are onto something, we need a more general model predicting and explaining their occurrence. Williamson (2013) does just this: he builds a model in epistemic logic which predicts the existence of Gettier cases, and explains when and why they occur. Williamson (rightly) holds that this is a substantial advance upon giving a counterexample to K = JTB, one which allows us to accept Gettier cases without reservations, since the model allows us to explain their occurrence and extent. Permissivists need to supply a similar model predicting and explaining the occurrence of permissive cases, rather than simply identifying putative counterexamples to uniqueness. In the absence of such a model, there are no assurances that the putative counterexamples are onto something, especially against the backdrop of powerful arguments in favor of uniqueness (of the sort which, I take it, the JTB analysis of knowledge never enjoyed).

These arguments, though similar in form to mine, in fact make very different points than I make.[15] Each of the three arguments identifies a distinct problem with the permissivist's commitment to a range of permissible credences. Horowitz's argument shows that permissivism undermines the most promising justification of the value of rationality: an accuracy first justification. Schultheis's argument shows that our imperfect ability to know where the edges of the permissible zone lie shows that permissivism is false. My argument shows that permissivists need to explain the scope of permissive cases, and that they cannot do so. These arguments complement each other quite well, and together the three constitute a very compelling case against the view that there is ever a range of permissible credences. I will have an occasion to further discuss the relation between my argument and those of Schultheis and Horowitz in section 4.

3. Answering the explanatory demands

The only kind of permissivism that has a principled answer to these explanatory demands is subjective Bayesianism. Bayesians endorse only two requirements on the nature of rational belief. First, one must have probabilistically coherent credences at each time. Second, one must update those credences by conditionalization on one's evidence. This epistemology is radically permissive. The only impermissive cases are tautologies and contradictions, for which the only permitted credence is 1 and 0, respectively. Other than that, every case is permissive. Bayesians explain the scope of epistemic permission in the following way: only full confidence is permitted to tautologies and zero confidence to contradictions because having full or zero confidence in a proposition expresses certainty that one will never gain any evidence that (dis)confirms the proposition. Tautologies and contradictions are the only propositions for which this confidence is warranted.

For contingent propositions, any credence is permissible, in the following sense. There are indeed formal constraints placed on the credences a subjective Bayesian might have in contingent propositions, but all of these constraints derive from pressures of coherence, be it coherence with the other attitudes the agent has at a time, or coherence with attitudes the agent had at an earlier time (the agent's present attitudes must be derived from her past attitudes by conditionalizing on the evidence she received in the meantime). Thus, for a single agent, it is not the case that any credence in any contingent proposition is permitted – the agent's attitudes must hang together in a specified way. But there are no substantive constraints on agents' credences in contingent propositions. This means that for any contingent proposition, body of evidence and credence, there is a subjective Bayesian agent who permissibly has that credence to that proposition given that body of evidence. So, every body of evidence paired with a contingent proposition is a permissive case. The scope of inter-personal permissions for Bayesians is indeed wide, even if the scope of intra-personal permissions is narrower.

I concede that subjective Bayesians can answer the explanatory demands that I have identified. This is because they are extreme permissivists: they think that almost all cases are permissive, and think that almost all attitudes are permitted in those permissive

---

[15]My argument is more comprehensive, in the sense that it identifies problems not only with the range of permissible attitudes within a permissive case, but also challenges the permissivist to explain the range of permission across the cases: to explain why some bodies of evidence are permissive with regard to some propositions, and others are not. There is no analogue of Schultheis's or Horowitz's argument in this broader case.

cases. And they have principled explanations of the few restrictions they place on epistemic permission, all having to do with the nature of logical truths and contradictions. But their overall epistemological outlook is highly implausible, given that they countenance no substantive constraints on rational belief.[16] Counter-inductivists and grue-projectors are counted as rational by subjective Bayesians, if they are coherent, since subjective Bayesians place no substantive constraints on our prior conditional or unconditional probabilities. According to subjective Bayesianism, someone with your same evidence could rationally have very high confidence that the moon is made of cheese, if they have suitably different priors. This extreme permissivism about the conditional and unconditional priors is very implausible. Thus, our task should be to find a *plausible* epistemological outlook which can answer the explanatory demands facing permissivism.

Schoenfield (2014) and Kelly (2013) offer approaches that might be able to do the job. According to Schoenfield, there are many different *epistemic standards* which one might reasonably adopt. She notes that epistemic standards can be conceived of in a variety of ways, but I think that it is most reasonable to think of them as different ways of balancing the many theoretical desiderata constitutive of good reasoning, like simplicity, strength, explanatory unification, non-arbitrariness, etc. Put this way, it is very similar to an older view of David Lewis's. Lewis held that the laws of nature were determined by the system of propositions which best weighed simplicity and strength, but that there is no single objectively privileged way to weigh these two desiderata. Instead, different theorists may identify different systems as best, depending on the way they weigh competing theoretical virtues.[17] Likewise, for Schoenfield, different inquirers may permissibly have different attitudes on the basis of the same evidence, but only if their standards differ. So we have a list of theoretical virtues, and a set of permissible ways to weigh those virtues against each other, and these are the possible epistemic standards. Two people with the same evidence but different standards could rationally have different attitudes to a proposition.[18]

Alternatively, Kelly (2013) suggests thinking of belief formation in terms of balancing the two Jamesian goals: believing true propositions and avoiding belief in false propositions. If these are our two goals in epistemic life, we may form beliefs differently depending on how we weigh the goals: someone who gives maximal weight to forming true beliefs and no weight to avoiding false belief will believe every proposition, and someone who gives the opposite weights will suspend judgement on every proposition. Though these two extremal weightings are impermissible, there are a variety of permissible intermediate weightings. People may permissibly weigh these two goals differently, and if this is the case, it could rationalize the fact that they permissibly have different attitudes on the basis of the same evidence.

Both Schoenfield's and Kelly's views are versions of interpersonal permissivism: once you've fixed the standards for a person, there is only one attitude that that person is permitted to have, in light of those standards and the person's evidence. But there

---

[16]An intuition I share with Horowitz (2014).

[17]Lewis (1994: 478–80). Lewis hoped that the world is "kind" so that every permissible way of weighing the theoretical desiderata would each designate the same system as best. If this were not the case, then there would be no determinate fact about what the laws are.

[18]See also Schoenfield (2019). Schoenfield (2022) argues that doubt, like one's epistemic standards, can be the source of permissivism, though whether we defer to the perspective of doubt once we take it up is not up to us. I think this view suffers from the same difficulties as standards permissivism (which I discuss below).

are a wide variety of permissible standards that one could have, and the fact that others have these standards and disagree with one although they share one's evidence does not imply that either party is irrational. In this way, Schoenfield and Kelly do not have to answer the explanatory demands concerning the range of permissive cases and how permissive the cases are *given a set of standards*, for there are no permissive cases when we relativize to a set of standards.

However, an exactly analogous explanatory demand faces Schoenfield and Kelly at the level of standards. What determines which standards are permissible and which are not? We need to explain the range of permissibility at the level of standards now, and this seems like a problem which is just as intractable. We can ask questions like: why is one way of weighing simplicity and strength permissible, but one that gives slightly more weight to strength impermissible? Or: why is one way of weighing the Jamesian goals permissible, but one that gives slightly more weight to believing truths impermissible. We will still have to draw a sharp, arbitrary line which needs explaining, this time between the permissible and impermissible standards rather than between permissible and impermissible cases or permissible and impermissible attitudes within a case. This explanatory problem at the level of standards is no easier to solve than the explanatory problem at the level of first-order permissive cases, because they are structurally isomorphic. Let a permissive* case be a case in which there is some set of permissible standards which rationalizes one's taking doxastic attitude D to proposition P given evidence E, and there is some set of permissible standards which rationalizes one's taking doxastic attitude D′ to proposition P given evidence E, where D = / = D′. Our new problem is to explain the scope of permissive* cases, and this is no easier than explaining the scope of permissive cases. Schoenfield's and Kelly's standards-based interpersonal permissivist approach merely relocates the explanatory demands, rather than answering them.

We have now seen that several different strategies for dealing with the explanatory demands either entailed an implausible epistemological outlook, or merely relocated, rather than answered, them. This is the dilemma for the permissivist: any view on which the explanatory demands do not arise will be a wildly implausible, extreme permissivism, like subjective Bayesianism. On the other hand, any view which uses some other factor to explain the scope of permission will give rise to new explanatory demands with regard to that factor. Since it looks like these explanatory demands cannot be answered, we shall now consider how they might be softened, by appealing to vagueness or ignorance.

4. Softening the explanatory demands

Two strategies are commonly appealed to in order to soften up explanatory demands like those facing the permissivist: identify a zone of vagueness or appeal to our ignorance of the relevant facts. When one is asked to explain the sharp line one draws between the things that are F and the things that are not-F, one can appeal to the vagueness of "F" to make the line less sharp, and hence, to lessen the force of the explanatory demand. Or, one could appeal to a class of cases for which we do not (or perhaps could not) know whether they are F or not to make the line between F and not-F *look* less sharp, and one could locate the tricky edge cases which generate the explanatory demands among these.[19] I shall consider these possibilities in turn, and argue that

---

[19]For epistemicists about vagueness, these strategies are the same.

neither will ultimately work: the permissivist has no way to soften up these explanatory demands.

4.1. Vagueness

Here is how the dialectic goes, put more concretely. Suppose that the range of permissible credences in P in a permissive case is [0.6, 0.8]. The uniquer asks why a credence of 0.8 is permitted but a credence of 0.800001 is not permitted. What is the relevant difference between a credence of 0.8 and 0.800001 such that the former is rational and the latter is irrational? A natural reply to this request for explanation is that the permissive range of credences does not have sharp borders. It is not the case that 0.8 is a determinately permissible credence but 0.800001 is a determinately impermissible credence. Rather, it is that the credences around 0.8 are such that it is indeterminate whether they are permissible or impermissible. This reply is attractive because it eliminates the sharp and seemingly arbitrary transition between the permissible and the impermissible credences which called out for explanation.

Here is another way to see this point. Consider a situation where we start out with no evidence pertaining to some proposition, P. Describe our evidential state at this time as E0. We then gradually gain evidence in small increments, E1, E2,…, En. Permissivists should say that E0 is a permissive case. Subjective Bayesians, for example, will simply allow an individual to have whatever credence in P in (0, 1) that they like before any evidence pertaining to P comes in. And we can pick a body of evidence for En such that the case in which someone has En as their evidence is clearly impermissive – it is clear that there is only one rational response to En. But each increment of evidence from Ei to Ei + 1 that one gains, we may suppose, is so minute that it makes no noticeable difference: if Ei is a permissive case then Ei + 1 is a permissive case. But if this is true, then En is also a permissive case. Contradiction. Permissivists could deny the inductive premise; they can do this in two ways. Permissivists can say that there is a sharp transition from permissive cases to non-permissive cases somewhere in the series, in which case they would face the acute explanatory demands that I have already identified. Alternatively, they could say that there is a vague frontier between the permissive and the impermissive cases. For cases in the penumbral area, some admissible precisifications of "permissive case" make the case come out as permissive, and some make it come out as impermissible. My argument in this part of the section will show that this way of denying the inductive premise will not work, either. But if the permissivist cannot deny the inductive premise, they must deny one of the two extremal premises. Denying that En is a permissive case would be to endorse nihilist permissivism, which I have already claimed is implausible. Denying that E0 is a permissive case would be to endorse the uniqueness thesis.[20]

Here is my argument in a reduced form:

P1. If permissivism is true, then in many permissive cases there are attitudes such that it is indeterminate whether they are permitted or not.

---

[20]It is important to note, in advance, that my argument does not generalize to show that the second way of denying the inductive premise of a sorites argument is not available, in general. (If my argument applied in the more general context, it would show that there were no heaps, or that one grain of sand constitutes a heap, an absurdity.) I believe that this move is perfectly legitimate in the case of non-deontic predicates for which a sorites argument can be set up. My argument is unique to deontic predicates like "permissive." Thanks to a reviewer at *Episteme* for encouraging me to clarify my argument in this respect.

P2. But there is no case such that it is indeterminate whether an attitude is permitted.

C. Hence, permissivism is not true.

I have just argued for P1: since permissivists cannot answer the explanatory demands that they face, they must instead weaken them by appealing to vagueness. I will now defend P2, by arguing that the notion of a borderline case of "permissive case" is incoherent. Indeed, my arguments support the stronger claim that the notion of anything's being indeterminately permissible is incoherent.

Here is one principle that I will rely on throughout my argument:

General Dominance: if doing A is determinately permissible, and doing B is indeterminately permissible, then all things considered, one should do A.

This dominance principle has a straightforward motivation: the fact that A is clearly permissible but it is not clear that B is permissible is something to be said in favor of A, it breaks the tie between A and B in favor of A. When deciding between which of two permissible options to perform, it takes very little to break the tie, since many things can serve as tie-breaking reasons. In our case, the fact that one of the options is not determinately permissible is more than enough to break the tie.

Similar dominance principles are standardly used in other normative domains such as in the theory of rational choice and in ethics. Debates about moral uncertainty, for example, often begin with a thought like the following: if both doing A and doing B will lead to outcomes with the same goodness, but doing B violates someone's rights and doing A does not, then even if you are highly confident in consequentialism – indeed, even if consequentialism is true – it is more appropriate, all things considered, to do A.[21] The possibility that deontology might be true is reason enough to break the tie between A and B in favor of A. In this context, "might" is an epistemic modality, but it seems that the motivations for the dominance reasoning apply just the same if we interpret the "might" in vagueness-theoretic terms instead. A is to be chosen instead of B in the case of moral uncertainty because the fact that some potential resolutions of our uncertainty hold that A is preferred to B, while others are indifferent between A and B. The fact that some resolutions of our uncertainty prefer A to B has the normative force to break a tie *even if* the correct resolution of our uncertainty leaves us indifferent between A and B. Likewise, A is to be chosen instead of B in the case of vagueness because the fact that some potential resolutions of the vagueness of "permissible" render A permissible and B impermissible, while others render both A and B permissible. The fact that some resolutions of the vagueness of "permissible" prefer A to B has the normative force to break a tie between A and B.[22] My claim here is simply that the

---

[21]MacAskill *et al.* (2020: 14–18), among others, use dominance reasoning of this kind to motivate their account of choice under moral uncertainty. On the theory of choice given moral vagueness in particular, see Dunaway (2017).

[22]This highlights a difference between the vagueness and uncertainty cases: there is no *privileged* resolution of vagueness like there is a privileged resolution of uncertainty, namely what is true (epistemicist theories of vagueness excepted). This means dominance reasoning, if it should be accepted in the case of ignorance, enjoys an even stronger justification in the case of vagueness.

analogy between uncertainty and vagueness helps us see that uncertainty and vagueness ought to play the same role in our reasoning – a role captured by the general dominance principle. That role is lexically posterior to first-order normative considerations, but that is all it takes to break a tie between options that otherwise enjoy parity.[23]

Indeed, the case for general dominance can be strengthened by considering it in the context of several particular theories of vagueness. I will consider three theories of vagueness: the degree theory, supervaluationism and epistemicism. These three theories correspond to three more specific forms of the dominance principle:

> *Degreed Dominance*: if the degree to which "A is permissible" is true is 1, but the degree to which "B is permissible" is true is less than 1, then all things considered, one should do A.
> *Supervaluative Dominance*: if all precisifications of "permissible" make "A is permissible" true, but not all precisifications of "permissible" make "B is permissible" true, then all things considered, one should do A.
> *Epistemic Dominance*: if one knows that doing A is permissible, but one does not know that doing B is permissible, then all things considered, one should do A.

These three specific dominance principles are simply the general dominance principle, combined with the respective theory of vagueness. I think each of these dominance principles is plausible, and consistent with the motivations of the respective theories. Schultheis ([2018](#)) defends a principle like epistemic dominance. Schultheis argues that if permissivism is true one could not know exactly where the ends of the permissible range of attitudes lie, and that one ought to have an attitude that one knows is permissible, rather than an attitude that one does not know to be permissible. A skeptic of epistemic dominance might retort that what matters is simply whether the attitude is permissible or not, not whether one knows that the attitude is permissible. But this retort claims that all that matters is *objective* rightness and wrongness, and that subjective rightness cannot even play a tie-breaking role. But this is not correct: even if objective rightness plays a lexically prior role in determining what to do, subjective rightness can play a tie-breaking role.[24] The epistemic dominance principle, recall, is the fundamental premise that motivates accounts of reasoning under moral uncertainty. The source of one's ignorance is different in the moral uncertainty case and the case of the vagueness of "permissible" (if epistemicism is true), but the source of the ignorance could not alter how it is rational to act given one's ignorance.

I will consider degree-theoretic and supervaluative dominance together, since they are both semantic theories of vagueness, and since supervaluationism simply collapses

---

[23]Here is a more speculative analogy. Some philosophers have claimed that the evaluative properties of something supervene on its modal properties, not merely on its occurrent properties. Kagan ([2019](#)), for instance, holds that modal persons, organisms that could have been persons (if things had gone differently), have a greater moral status than modal, non-persons. If the fact that something might have been a person should influence how we treat it (where the "might" is that of metaphysical modality), then perhaps the fact that an option might be impermissible should be a reason for us to choose another option when one is available (where the "might" is that of indeterminacy).

[24]Compare this to the account of norms in Williamson (forthcoming). A and B both satisfy the first-order norm of having a permissible attitude, but having attitude A (which one does not know to be permissible) violates a higher-order norm of having a disposition to satisfy first-order norms. Violating a higher-order norm and conforming to a first-order norm, I contend, is inferior to conforming to both a first- and a higher-order norm.

into the degree theory when there is a measure on the class of admissible precisifications of "permissible." The fundamental motivation for these principles remains the same as above: it takes very little to break a tie between two permissible options. Vagueness due to semantic indecision can play this role. Your conduct should be such that it is acceptable, no matter how we sharpen up our language. Scrupulous actors and thinkers should take all precautions to ensure that their conduct and attitudes are permissible. Having an attitude that would be impermissible if we sharpened up our usage in a particular way is not something a scrupulous thinker would do, if a safer attitude is available. In general, if two options are tied with respect to the primary criterion of evaluation, then anything to be said in favor of the first option that is not to be said in favor of the other is sufficient to cause a perfectly rational actor to choose the first option. Degree-theoretic and superevaluative dominance are valid principles of reasoning, then, if the fact that all sharpenings of our moral and epistemic language, and the practices that underlie them, support an option is something to be said in favor of that option, and I contend that it is.[25]

I need a second principle, in addition to general dominance. That principle is:

Obligation implies permission: If one is obligated to do A, then one is permitted to do A.

This principle is the best candidate for an analytic truth about deontic concepts, and as such it is a suitable foundation on which to build any deontic logic, so I will not defend it any further here.

Now that I have identified and defended my crucial assumptions, I will present my defense of premise 2. Suppose for *reductio* that C is a borderline case of "permissive case." If C were a determinate permissive case, then there would be two doxastic attitudes, A and B, such that one is permitted to have A (PA) and one is permitted to have B (PB). So (PA and PB) would be determinately true. But given that C is an indeterminate permissive case, it is indeterminate that (PA and PB). If it is indeterminate that (PA and PB), exactly one of the following must be true, on any logic of truth functions under indeterminacy:[26]

Horn 1: It is indeterminate that PA and determinate that PB.[27]
Horn 2: It is indeterminate that PA and it is indeterminate that PB.

This sets up a dilemma.
Here is horn 1. There are two ways for PA to be indeterminate:

---

[25]The argument I am giving here mirrors Broome's (1997) justification of what he calls the "collapsing principle" in a discussion of vagueness and value incommensurability. See also Carlson (2013) for a good critical discussion. Broome argues that if it is not false that X is better than Y and it is false that Y is better than X, then it is true that X is better than Y. Transposing this principle to our context: if it is not true that X is permissible, and it is true that Y is permissible, then it is false that X is permissible. I believe that Broome's arguments carry over to support this principle, which entails the superevaluative and degree-theoretic dominance principles.
[26]Recall that a conjunction cannot be indeterminate if neither of its conjuncts is indeterminate, so it is not a trilemma.
[27]Or it is indeterminate that PB and it is determinate that PA, but we can consider only the first case without loss of generality.

Way 1: A lies in the vague frontier between permissibility and obligatoriness.
Way 2: A lies in the vague frontier between permissibility and forbiddeness.

Consider way 1. Option A sits in between obligation and permission. Since obligation implies permission, option A in fact sits between permission and permission. How could it be that an option is indeterminately permissible if it is sandwiched between two determinately permissible options? I think this is impossible. I will illustrate this by assuming that epistemicism is true. On the first horn, option A is either permissible or obligatory, but which one it is simply is unknowable. But we can reason as follows: if A is obligatory, then it is permissible because obligation implies permission, but if A is permissible, then A is permissible, because any proposition implies itself. So either way, A is permissible. In other words, this option is not a borderline case of permissibility at all: all ways of resolving the indeterminacy imply that the option is permissible. Instead, the option is a borderline case between two different types of permissibility: the type that is implied by obligation and the type that is not.

Now consider way 2. This is straightforwardly ruled out by the dominance principle I identified above. Since there is an option that is determinately permissible, B, the dominance principle rules out the indeterminately permissible option A. A is all things considered a worse option than B, because it *might* be impermissible. Regardless of whether this "might" is spelled out in epistemic, supervaluationist or degree theoretic terms, I have argued that an option such that it is not the case that it might be impermissible should be preferred to one that might be impermissible.

Horn 2 is where things get interesting. It is indeterminate that A is permissible *and* it is indeterminate that B is permissible. There are three ways for this to occur:

Case 1: A and B both lie in the vague frontier between permission and obligation.
Case 2: A and B both lie in the vague frontier between permission and the forbidden.
Case 3: A lies in the vague frontier between permission and the forbidden and B lies in the vague frontier between permission and obligation.[28]

In case 3, we can employ reasoning that we used on the first horn of the dilemma. Since B lies in the vague frontier between obligation and permission, we can conclude that it is not really an indeterminate case of permission after all, since on every way of resolving the indeterminacy, the option is permissible, given that obligation implies permission. Then, the dominance principle rules out option A, which is indeterminately permitted, because there exists another option that is determinately permitted.

Cases 1 and 2 are instances of what I shall call *quasi-dilemmas*. Classical dilemmas are cases in which you are determinately obligated to ϕ and you are determinately obligated to ψ, but you cannot both ϕ and ψ, or cases in which you are determinately forbidden to ϕ and you are determinately forbidden to ψ, but you cannot avoid either ϕ-ing or ψ-ing. Quasi-dilemmas are cases in which you are indeterminately obligated to ϕ and you are indeterminately obligated to ψ, but you cannot both ϕ and ψ, or cases in which you are indeterminately forbidden to ϕ and you are indeterminately forbidden to ψ, but you cannot avoid either ϕ-ing or ψ-ing.[29]

---

[28]Or vice versa, but we can consider only one case without loss of generality.

[29]When I say that an act is "indeterminately forbidden" or "indeterminately obligatory" in this sentence I mean that it lies in the vague frontier between obligation and permission, or the vague frontier between permission and being forbidden. Interestingly, there is a third kind of dilemma, one in which ϕ is

I shall now argue against cases 1 and 2 as follows:[30]

P1. There are no dilemmas.
P2. If there are no dilemmas, then there are no quasi-dilemmas.
C. There are no quasi-dilemmas.

At the very least, it is dialectically inappropriate for permissivists to hold that there are dilemmas, and rely on this fact in arguments with uniquers. This is because epistemic dilemmas are permissive cases, given only the assumption that ought implies permission, a very weak assumption of standard deontic logic. In a dilemma where one is obligated to do A and obligated to do B (but can't do both), then one is permitted to do A and permitted to do B – a permissive case! The existence of epistemic dilemmas entails permissivism, but the non-existence of dilemmas is compatible with both uniqueness and permissivism, so that is what must be assumed in debates between permissivists and uniquers. Otherwise, permissivists will have begged the question. I have identified one feature of dilemmas which makes it dialectically inappropriate to assume that there are dilemmas in our present context, but there are many other arguments against epistemic dilemmas as well.[31] There are also many compelling arguments against dilemmas in the moral case – ones which could easily be repurposed for the epistemic case – that I cannot rehearse here.

It is pretty easy to see how arguments against dilemmas carry over straightforwardly as arguments against quasi-dilemmas. I regard quasi-dilemmas as just as bad as classical dilemmas: they are both cases in which it is not possible to fulfill all of your (determinate or indeterminate) obligations, or it is not possible to avoid everything that is (determinately or indeterminately) forbidden. The same kind of logical and methodological scruples most theorists have about dilemmas will also apply to quasi-dilemmas. My argument concerning obligation implies permission carries over particularly well. Indeterminate obligation implies permission, because if you are indeterminately obligated to do something, every way of resolving that indeterminacy entails that you are permitted to do it. But then, a quasi-dilemma would just be a permissive case, for if one is indeterminately obligated to do A, then one is permitted to do A, and if one is indeterminately obligated to do B, then one is permitted to do B, a permissive case. By positing quasi-dilemmas, the permissivist begs the question against the uniquer. We have now seen that it is dialectically inappropriate to assume, in a debate between permissivists and uniquers, that there are epistemic dilemmas or quasi-dilemmas, yet, this is where the permissivist's argument ended up.

This completes my *reductio* against the supposition that there is an indeterminate permissive case, and as such, allows me to draw the conclusion that permissivism is

---

indeterminately obligatory and ψ is determinately obligatory (*mutatis mutandis* for forbidden). However, I think that the dominance reasoning I have relied on in the text can be extended to a dominance principle concerning obligation, one that would show that one should, all things considered, perform the determinately obligatory option rather than the indeterminately obligatory option in this third kind of dilemma.

[30]I will lump together cases 1 and 2 in my discussion, because their logical structure is the same. If one is forbidden from doing A, and one is forbidden from doing B, then one is obligated not to do A, and one is obligated not to do B. Since A and B are exclusive attitudes, one cannot fulfill all of one's obligations, a dilemma.

[31]See Hughes (forthcoming) for critical discussions of many of these arguments.

not true.[32] I argued that there is no room for vagueness about deontic facts concerning what is permitted, forbidden or obligatory, and as such, that the very notion of an indeterminate permissive case is misguided.[33] If the permissible range of attitudes must have sharp borders, then, the permissivist faces a daunting explanatory demand that it was this paper's primary objective to identify.

4.2. Ignorance

Perhaps the permissivist could weaken the explanatory demands they face by holding that, though there is in fact a sharp cutoff between permissible and impermissible, we are ignorant of where this cut off lies. This is the reason why it appears to us as if there is no sharp cutoff between the permitted and the impermitted attitudes. We do not know, of the cases which are around the true sharp cutoff between the permitted and impermitted attitudes, whether they are permitted or not.[34]

This reply has some virtues. For example, it will tend to make sense to anyone with a conception of evidence according to which our evidence is not guaranteed to be first-personally accessible, or, perhaps even one according to which we can make mistakes in determining what our evidence supports. Anyone who holds the E = K thesis, coupled with a failure of positive introspection, like Williamson (2000), will have such a position.

But there are three problems with this reply. First, Schultheis (2018) argues that if you do not know where the endpoint of the permissive zone of credences is located, then having a credence on the edge of the permissive zone is not in fact permissible, because you should never have an attitude which is such that you are not sure whether it is permissible, when there is another attitude which you know to be permissible. But since the endpoint of the zone of permissible credences is not in fact permissible, there is no such permissible zone. That argument, if it goes through, would preclude the permissivist from appealing to our ignorance to explain why it seems like there is not a sharp transition between the permissive cases and the non-permissive cases, for doing so would falsify permissivism if Schultheis's reasoning is correct.

Second, any such account will compromise the initial motivation for permissivism itself. Here is why. If we explain the appearance that there are not sharp borders to the permissive zone by appealing to our ignorance of where those borders lie, we might as well explain (away) the intuitions behind permissivism itself in a similar manner. Such an explanation would hold that though it appears that in some cases a range of attitudes is rationally permitted, in fact, only one attitude is rationally permitted, but we do not know which of a number of attitudes this is. On this account, the zone of attitudes which appear to be permitted would be composed of attitudes such that we do not know whether the attitude is permitted or not, though in reality, only one

---

[32]Admittedly, there are a lot of moving parts in the argument, and it does not rely on a formal deontic logic with indeterminacy. To my knowledge, no such logic exists, and no doubt, the development of such a logic would improve our understanding of uniqueness and permissivism, as well as a number of other issues in normative inquiry.

[33]Matheson (2011: 371) also discusses the issue of vagueness as it relates to permissivism and uniqueness.

[34]Epistemicists about vagueness also think that there will be a sharp cutoff and that we will not know where it lies, but they hold that our ignorance comes from a special source – the fact that usage could have easily been different. See Schoenfield (2016) on this point. Here I will discuss the broader class of ignorance-based explanations.

attitude in that zone is permitted; our ignorance of which attitude this is gives rise to the appearance that there are many permissible attitudes. The moral here is that if we are going to explain the fuzziness of the edges of the permissible zone by appealing to our ignorance of where its borders lie, then there is going to be theoretical pressure coming from a concern for unification and simplicity to explain the whole appearance that there is a range of permitted attitudes in a similar manner – by appealing to our ignorance of which is the uniquely rational attitude. If we can offer a unified explanation of the appearance that many attitudes are permissible by appealing to our ignorance, then we should do so, rather than offering a disunified explanation of the appearances by explaining the apparent fuzziness of the edges of the permissive zone by appealing to our ignorance and explaining its existence as owing to the fact that there really are multiple permissible options. Hence, uniqueness + ignorance about which attitude is permitted is a simpler explanation than permissivism + ignorance about the location of the endpoints of the permissible zone, and as such the former should be adopted in preference to the latter unless there are compelling competing considerations.

Third, by replying in this way, the permissivist is still committed to the existence of sharp transitions between permissive and non-permissive cases and between permitted and impermitted attitudes in a permissive case, and as such, still carries an explanatory burden. Appealing to our ignorance shields the phenomena from explanatory demands only by obscuring them; it does not ultimately answer the explanatory demands. The sharp transition between permitted and impermitted attitudes, or between permissive and impermissive cases, wherever it is, will be arbitrary, whether we know about it or not. The permissivist still needs to offer us an epistemological theory about what explains the extent of epistemic permission within and across the cases – whether or not they obscure the sharp transitions by claiming that we do not know where they lie.[35]

This completes my argument that permissivists cannot appeal to indeterminacy or ignorance to soften the explanatory burdens that they must shoulder. The permissivist must still explain which cases are permissive and why, and how permissive they are, and

---

[35]Incidentally, these considerations highlight one respect in which my argument is not susceptible to an objection that faces Schultheis's and Horowitz's. Both of those arguments are vulnerable to a reply like the following: when one is in a permissive case, one must believe that one's own credence is in the middle of the permissible range of credences. And when one believes that one's credence is in the middle of the range of permissible credences, Schultheis's and Horowitz's arguments do not apply. Call this view "skewed permissivism." This blocks Horowitz's argument because all credences one regards as permissible are then more expectedly accurate than any impermissible credence, and it blocks Schultheis's argument because one will never have an attitude which one takes to be on the edge of the permissible zone. Schultheis and Horowitz both consider this reply, and criticize it because it commits the skewed permissivist to knowing in advance that they will often have false beliefs about which attitudes are rational; this will be the case whenever they do not have credences which are actually at the center of the range of permissible attitudes. There is also a stronger objection, which Smith (2020) defends and calls "unacknowledged permissivism," according to which one can never even know that one is in a permissive case. But a similar reply works here: it seems odd that the permissivist is committed now to unknowable facts about which attitudes are rationally permissible, and about which cases are permissive. I think that Schultheis's and Horowitz's replies to the skewed and unacknowledged permissivist are convincing. However, it is worth noting that my argument still goes through even if these replies fail and skewed or unacknowledged permissivism can be made to work. This is because I am making a metaphysical point: the permissivist must explain the scope of the facts about permission among cases and within a case, whether or not an agent knows these facts. By contrast, the arguments of Schultheis and Horowitz depend on an agent's knowledge of some facts about the scope of permission, which is what the unacknowledged and skewed permissivists require to get their objections going.

why. Adding in a dose of indeterminacy or ignorance will not help – indeed, it hurts because it further exposes the vulnerabilities of permissivism.

5. Conclusion

In this paper, my main contribution has been to develop a family of explanatory demands facing any version of epistemic permissivism. Permissivists usually argue for their view by trying to identify a counterexample to the uniqueness thesis. Yet, permissivists have not recognized that by offering such counterexamples, they incur an explanatory burden. They must explain why some cases are permissive and others are not, and they must explain for each permissive case, how permissive it is. The reason this explanatory demand will prove difficult or impossible to fulfill is that the lines permissivists must draw between permitted and impermitted cases and attitudes are arbitrary. Only subjective Bayesians have principled answers to these explanatory demands, but their epistemological outlook is so implausible that even most permissivists are unwilling to accept it. And when permissivists have attempted to offer an explanation of why permissive cases arise, they have merely relocated, rather than answered, the explanatory demands. Finally, we saw that attempts to soften these explanatory demands by appealing to ignorance or vagueness failed.

In short, permissivists must tell us what determines whether a case is permissive, and how permissive it is. In the absence of such an explanation, uniquers are more than justified in their practice of explaining away the intuitions behind putative counterexamples to uniqueness. The fact that permissivists are not equipped to answer this explanatory demand is indeed compelling, if not decisive, evidence against their view.[36]

References

Broome J. (1997). 'Is Incommensurability Vagueness?' In R. Chang (ed.), *Incommensurability, Incomparability and Practical Reason*, pp. 1–17. Cambridge, MA: Harvard University Press.

Carlson E. (2013). 'Vagueness, Incomparability, and the Collapsing Principle.' *Ethical Theory and Moral Practice* **16**(3), 449–63. doi: 10.1017/s0953820804000548.

Castro C. and Hart C. (2019). 'The Imprecise Impermissivist's Dilemma.' *Synthese* **196**(4), 1623–40. https://doi.org/10.1007/s11229-017-1530-9.

Dunaway B. (2017). 'Ethical Vagueness and Practical Reasoning.' *Philosophical Quarterly* **67**(266), 38–60. https://doi.org/10.1093/pq/pqw038.

Feldman R. (2006). 'Epistemological Puzzles about Disagreement.' In S. Hetherington (ed.), *Epistemology Futures*, pp. 216–36. Oxford: Oxford University Press.

Goldman A. (2010). 'Epistemic Relativism and Reasonable Disagreement.' In R. Feldman and T. Warfield (eds), *Disagreement*. pp 187–215, New York, NY: Oxford University Press.

Greco D. and Hedden B. (2016). 'Uniqueness and Metaepistemology.' *The Journal of Philosophy* **113**(8), 365–95. https://doi.org/10.5840/jphil2016113825.

Horowitz S. (2014). 'Immoderately Rational.' *Philosophical Studies* **167**(1), 41–56. https://doi.org/10.1007/s11098-013-0231-6.

Hughes N. (ed.) (forthcoming). *Epistemic Dilemmas*. Oxford: Oxford University Press.

Jackson E. (2021). 'A Defense of Intrapersonal Belief Permissivism.' *Episteme* **18**(2), 313–27. https://doi.org/10.1017/epi.2019.19.

Kagan S. (2019). *How to Count Animals, More or Less*. Oxford, UK: Oxford University Press.

Kelly T. (2013). 'Evidence Can Be Permissive.' In M. Steup and J. Turri (eds), *Contemporary Debates in Epistemology*, p. 298. Hoboken, NJ: Blackwell.

Lewis D. (1994). 'Humean Supervenience Debugged.' *Mind* **103**(412), 473–90.

MacAskill W., Bykvist K. and Ord T. (2020). *Moral Uncertainty*. New York: Oxford University Press.

Matheson J. (2011). 'The Case for Rational Uniqueness.' *Logic and Episteme* **2**(3), 359–73.

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

Schoenfield M. (2014). 'Permission to Believe: Why Permissivism Is True and What It Tells Us about Irrelevant Influences on Belief.' *Noûs* **48**(2), 193–218. https://doi.org/10.1111/nous.12006.

Schoenfield M. (2016). 'Moral Vagueness Is Ontic Vagueness.' *Ethics* **126**(2), 257–82. https://doi.org/10.1086/683541.

Schoenfield M. (2019). 'Permissivism and the Value of Rationality: A Challenge to the Uniqueness Thesis.' *Philosophy and Phenomenological Research* **99**(2), 286–97. https://doi.org/10.1111/phpr.12490.

Schoenfield M. (2022). 'XII – Deferring to Doubt.' *Proceedings of the Aristotelian Society* **122**(3), 269–90. https://doi.org/10.1093/arisoc/aoac011.

Schultheis G. (2018). 'Living on the Edge: Against Epistemic Permissivism.' *Mind* **127**(507), 863–79. https://doi.org/10.1093/mind/fzw065.

Smith J. (2020). 'Unacknowledged Permissivism.' *Pacific Philosophical Quarterly* **101**(1), 158–83. https://doi.org/10.1111/papq.12299.

Titelbaum M.G. and Kopec M. (Manuscript). 'Plausible Permissivism.'

White R. (2005). 'Epistemic Permissiveness.' *Philosophical Perspectives* **19**(1), 445–59. https://doi.org/10.1111/j.1520-8583.2005.00069.x.

Williamson T. (2000). *Knowledge and Its Limits*. Oxford, UK: Oxford University Press.

Williamson T. (2013). 'Gettier Cases in Epistemic Logic.' *Inquiry* **56**(1), 1–14.

Williamson T. (Forthcoming). 'Justifications, Excuses, and Sceptical Scenarios.' In F. Dorsch and J. Dutant (eds), *The New Evil Demon*. Oxford: Oxford University Press.

Paul Forrester is a PhD candidate in philosophy at Yale University. His research mainly involves ethics, epistemology, and metaethics. Email: m.paul.forrester@gmail.com