Permissive Situations and Direct Doxastic Control[*]

BLAKE ROEBER
University of Notre Dame

Abstract: According to what I will call 'the disanalogy thesis,' beliefs differ from actions in at least the following important way: while cognitively healthy people often exhibit direct control over their actions, there is no possible scenario where a cognitively healthy person exhibits direct control over her beliefs. Recent arguments against the disanalogy thesis maintain that, if you find yourself in what I will call a 'permissive situation' with respect to p, then you can have direct control over whether you believe p, and do so without manifesting any cognitive defect. These arguments focus primarily on the idea that we can have direct doxastic control in permissive situations, but they provide insufficient reason for thinking that permissive situations are actually possible, since they pay inadequate attention to the following worries: permissive situations seem inconsistent with the uniqueness thesis, permissive situations seem inconsistent with natural thoughts about epistemic akrasia, and vagueness threatens even if we push these worries aside. In this paper I argue that, on the understanding of permissive situations that is most useful for evaluating the disanalogy thesis, permissive situations clearly *are* possible.

Epistemologists have grown increasingly interested in the question how epistemic rationality compares to practical rationality (Berker 2013, Cohen 2016, Rinard 2017, *etc.*), but epistemologists have been asking the more general question how belief relates to action for a very long time—for at least as long as they've wondered whether, and to what extent, we can control our beliefs.

According to the currently dominant view, beliefs differ from actions in at least this way: while we often have direct control over our actions, we never have direct control over our beliefs. On this view, just as I might cause myself to blush by thinking about something embarrassing, I might cause myself to believe that my lights are on by looking at my lights while turning them on (Feldman 2001). But on this view, I can't form the belief that the lights are on, or any other belief, by simply *deciding* to form it, the way I can (for example) raise my arm by simply deciding to raise it. On this view, if direct control over our beliefs isn't fully conceptually impossible, it's at least impossible for properly functioning people like you and me. Perhaps Bennett's Credamites can do it (1990), but *we* can't do it without getting ourselves into a defective cognitive state like theirs. Thus, while cognitively healthy people often exhibit direct control over their *actions*, there is no possible scenario where a cognitively healthy human being exhibits direct control over her *beliefs*.

Call this thesis about the relationship between belief and action the 'disanalogy thesis.' As Kurt Sylvan notes (2016), Joseph Raz (1999), Carl Ginet (2001), Keith Frankish (2007), Philip Nickel (2010), and Conor McHugh (2013) all reject the disanalogy thesis, and they all reject it because they all think that, if a person found herself in a situation where she's rationally permitted to believe some proposition and simultaneously rationally permitted to suspend judgment on that proposition, she can have direct control over whether she believes it *without* exhibiting any kind of cognitive defect. I think these authors are on to something. But it's not obvious that there *are* situations where we're rationally permitted to believe some proposition and simultaneously

---

rationally permitted to suspend judgment on that proposition, since these situations seem inconsistent with the arguments for uniqueness in the permissivism/impermissivism debate, Kurt Sylvan (2016) has explicitly argued that epistemic akrasia renders these situations impossible, and vagueness threatens even if we push these worries aside. On the best way to understand permissive situations, however, permissive situations survive literally every argument for the uniqueness thesis, Sylvan's arguments give us insufficient reason to think that akrasia threatens permissive situations, and the complications that arise from vagueness only *support* the general idea behind permissive situations. Or so I shall argue in this paper.

1. The Argument from Permissive Situations

From here forward, let's say that doxastic attitude A toward p fits body of total evidence E just in case E provides at least as much support for A as it provides for any alternative doxastic attitude toward p, let's say that doxastic attitude A toward p *uniquely* fits E just in case E provides *more* support for A than it provides for any alternative doxastic attitude toward p, and let's say that doxastic attitude A toward p *does not* fit E just in case E provides more support for some *alternative* doxastic attitude toward p than it provides for A. When does E provide at least as much support for A as it provides for any alternative attitude, then, or more support for A than it provides for any alternative attitude? By 'support,' I mean something half normative and something half psychological, in line with the proper-functionalist views defended by Plantinga (1993), Bergmann (2006), and others.

   Right now, my total evidence compels me to believe that I'm sitting at my desk, and it compels me to suspend judgment on the proposition that the number of stars is even. Even if I wanted to, I couldn't form the belief that the number of stars is even, or shake the belief that I'm sitting at my desk. Right now, my total evidence also compels me to believe that I had a bagel for breakfast, but this compulsion isn't as strong as my compulsion to believe that I'm presently sitting at my desk, since it's not as obvious to me that I had a bagel for breakfast as it is that I'm sitting at my desk. Similarly, my total evidence compels me to suspend judgment on the proposition that it will be sunny tomorrow, but this compulsion isn't as strong as my compulsion to suspend judgment on the proposition that the number of stars is even, since it's more plausible that it will be sunny tomorrow than it is that it will be be cloudy tomorrow, but not more plausible that the number of stars is even than it is that the number of stars is odd. For different propositions, doxastic compulsions like these push with different strengths toward different attitudes. Now, presumably, I'm not suffering any cognitive defect in experiencing these compulsions. (I'm not, for example, sitting on the beach in Bali, watching the sun set, while feeling fully compelled to believing that I'm sitting at my desk.) Presumably, *any* cognitively healthy person with exactly my total evidence would experience these compulsions, or at least fairly similar ones, if she were considering the same questions in roughly the same circumstances.

   When I talk about evidential support, I'm not talking about the probability of a proposition conditional on the relevant evidence—at least not directly. Instead, I'm talking about the level of compulsion that a cognitively healthy person who had that evidence would feel if she were considering the relevant question. More carefully, when I say that E provides more support for A

than it provides for any alternative doxastic attitude, I mean (to a first approximation) that a cognitively healthy human being who was considering the relevant question and who had E as her total evidence would feel a compulsion toward A that was *stronger* than any compulsion she felt toward any alternative doxastic attitude. And similarly, when I say that E provides at least as much support for A as it provides for any alternative doxastic attitude, I mean that a cognitively healthy human being who was considering the relevant question and who had E as her total evidence would feel a compulsion toward A that was *at least as strong* as any compulsion she felt toward any alternative doxastic attitude. So, on the stipulative definition of 'fit' that I will employ throughout, doxastic attitude A toward p fits E just in case a cognitively healthy human being who was considering the question whether p and who had E as her total evidence would feel a compulsion toward A that was at least as strong as any compulsion she felt toward any alternative doxastic attitude toward p.

With this definition of 'fit' in hand, I can introduce two more terms of art: 'permissive situations' and 'impermissive evidentialism.' By the former, I mean situations where belief in some proposition and suspension of judgment with respect to that proposition *both* fit a given body of total evidence. By the latter, I mean the thesis that, for every proposition p and body of total evidence E, some doxastic attitude toward p uniquely fits E. Impermissive evidentialism entails that permissive situations aren't possible, since it entails that a given attitude toward p fits E just in case it *uniquely* fits E. Thus, if impermissive evidentialism is true, there is no possible scenario where belief in some proposition and suspension of judgment with respect to that proposition both fit a given body of total evidence.

As I noted above, Raz, Ginet, Frankish, Nickel, and McHugh all think that, when a person finds herself in a situation where believing some proposition and suspending judgment on that proposition are both rationally permissible, she can have a significant form of direct control over whether she believes the proposition in question. As L.J. Cohen notes in the opening paragraph of his *Blackwell Companion to Epistemology* entry on rationality, however, "there are at least nine types of rationality" (2010: 663); and as Plantinga (1993), Worsnip (2015), and others make clear, there are more than just the nine types enumerated by Cohen. So, when we talk about rationality, it's important to state clearly what we mean. On something like Plantinga's proper-functionalist conception of rationality (ibid: 136-7), we can say that a person's belief is rationally permissible just in case it fits her total evidence, in the sense of 'fit' defined above. I'm not sure that Raz, Ginet, Frankish, Nickel, and McHugh are thinking about rationality or related epistemic properties this way, and I suspect that some of them are not. But I think this is the conception of rationality that gives us the strongest version of their respective arguments, since it's the conception of rationality best-suited for evaluating the idea that there is no possible scenario where a *cognitively healthy* person exhibits direct control over one of her beliefs. So, as I will read them in this paper, Raz, Ginet, Frankish, Nickel, and McHugh all endorse something like the following argument.

(1)   Permissive situations are possible.

(2)   If permissive situations are possible, then a cognitively healthy person might have as much control over one of her beliefs as she has over any of her actions.

∴   (3)   The disanalogy thesis is false.

I accept the conclusion of this argument, but I won't defend that conclusion here, since (2) requires a paper of its own. I defend (2) in my paper "Evidence, Judgment, and Belief at Will," and Raz, Ginet, Frankish, Nickel, and McHugh have given their own arguments for (2).[1] In this paper, I defend (1).

2. Doxastic *versus* Credal Impermissive Evidentialism

In §1, I defined 'impermissive evidentialism' as the view that a given body of total evidence must always provide more support for some doxastic attitude toward a proposition than it provides for any alternative doxastic attitude toward that proposition, but I didn't say what I meant by 'doxastic attitudes.' I meant just belief, disbelief, and suspension of judgment, and this is how I will use the term throughout. On many theories of belief, however, you and I might have different doxastic attitudes toward a proposition without having different *credences* in that proposition.[2] With this in mind, compare the following versions of impermissive evidentialism.

> **Doxastic Impermissive Evidentialism (DIE)**: For every proposition p and body of total evidence E, some doxastic attitude toward p uniquely fits E.

> **Credal Impermissive Evidentialism (CIE)**: For every proposition p and body of total evidence E, some credence in p uniquely fits E.

DIE and CIE aren't equivalent theses. To see why, consider views of belief like the one that Jacob Ross and Mark Schroeder (2014) call 'pragmatic credal reductivism,' or 'PCR' for short.[3] According to PCR, believing p consists in having a sufficiently high credence in p, where your credence in p counts as sufficiently high just in case it is high enough to ensure that you are willing to act as if p in some specified set of circumstances (p. 263). On this view, if your credence in p fits your total evidence and this credence is sufficiently high to count as a belief in p, then you believe p and your belief in p fits your total evidence. Similarly, on this view, if your credence in p fits your total evidence and this credence is sufficiently low to count as some alternative doxastic attitude toward p (a suspension of judgment with respect to p or a disbelief in p), then you have this alternative doxastic attitude toward p and this alternative doxastic attitude toward p fits your total evidence. So now assume that CIE is true, assume that (say) 0.85 is the credence in p that uniquely fits our total evidence, and assume that, in virtue of some non-evidential difference in our values, this credence is high enough to ensure that you are willing to act as if p in the relevant set of circumstances but not high enough to ensure that I am willing to act as if p in these

---

[1] See the books and articles mentioned above. As Sylvan points out (2016: 1637), Ginet, Frankish, and Nickel (on the one hand) and Raz and McHugh (on the other) offer importantly different reasons for thinking that (2) is true. Steup 2017 (§§3-5 in particular) is also relevant in this context.

[2] See the theories and arguments in Ganson 2008, Sosa 2010, Weatherson 2012, Friedman 2013, Buchak 2014, Ross and Schroeder 2014, Locke 2014, and Staffel 2016, to name just a few.

[3] According to Ross and Schroeder (p. 260), Ganson 2008, Fantl and McGrath 2009, and Weatherson 2012 all defend PCR.

circumstances (p. 278). Then, if PCR is true, your having credence 0.85 in p would amount to your believing p while my having credence 0.85 in p would amount to my having some alternative doxastic attitude toward p, and, moreover, your belief in p would fit your total evidence while my alternative attitude toward p would fit my total evidence, even though (by hypothesis) we have exactly the *same* total evidence. But this conclusion conflicts with DIE, since it denies that, for every proposition p and body of total evidence E, some doxastic attitude toward p uniquely fits E. So, if PCR is true, CIE and DIE aren't just different theses, they're potentially *conflicting* theses.[4] Of course, it's an open question how beliefs relate to credences, and PCR is a controversial view. But PCR suffices to show that we shouldn't treat CIE and DIE as equivalent views.

In §3, I will argue that permissive situations are possible, and that DIE is therefore false. And since I think that permissive situations are possible even if CIE is true, I will assume throughout that CIE *is* true. But I won't rely on PCR or any other specific theory of belief in my argument for permissive situations. Instead, I will simply assume that the *naïve* threshold view of belief is false.[5]

3. The Argument *for* Permissive Situations

Why think that permissive situations are possible? The short answer is that they *seem* possible and nobody has produced any good reason for thinking otherwise. In this section, I will say why permissive situations seem possible. In §§4-6, I'll explain why nobody has produced any good reason for thinking that they aren't possible.[6]

Suppose you want to know whether p is true, and consider the various evidential situations in which you might find yourself. In the best-case scenario (strictly epistemically speaking), your total evidence decisively answers the question whether p is true, either by making it certain that p is true, or by making it certain that p is false. In the worst-case scenario (strictly epistemically speaking), your total evidence doesn't give you the slightest hint whether p is true, and a blind guess whether p is true is the best you can do. Plausibly, suspension is the only attitude that fits your total evidence in the latter scenario, while either belief or disbelief is the only attitude that fits your total evidence in the former scenario (depending on which direction your evidence points). So, plausibly, DIE gets the right results in both the best-case scenario and the worst-case scenario.

---

[4] Katherine Rubin (2015) defends a related point, and Dustin Locke's discussion of the people he calls Spike and Mike (2014: 50-51) raises similar issues.

[5] By 'the naïve threshold view of belief,' I mean the view that there is some fixed, context invariant value of 'x' (0.75 or something) for which believing a proposition *just is* having a credence in that proposition at least as high as x. I know of no philosopher who explicitly accepts the naïve threshold view of belief, and the more sophisticated threshold views that we find in the literature—Ganson's (2008), Weatherson's (2012), Clarke's (2013), Leitgeb's (2014), *etc.*—all entail that the naïve threshold view of belief is false. (And of course, every *non*-threshold view of belief also entails that the naïve threshold view of belief is false.)

[6] The view that I sketch in this section is similar to the view presupposed in Nickel's argument that belief and suspension are both rationally permissible in what he calls his "target cases" (2010: 315-16). As Sylvan points out, however, it's easy to think that in Nickel's cases you're only managing to suspend judgment by manipulating yourself into thinking that you lack sufficient evidence for belief (2016: 1645). As will be clear in §5, this is one of several important differences between Nickel's view and the view that I sketch here.

But even supposing that DIE gets the right results in both the best-case scenario and the worst-case scenario, what should we think happens if you start in the worst-case scenario and go by gradual improvements in your evidence to the best-case scenario? What should we think happens if you start out in a situation where suspension is the only attitude that fits your total evidence and you slowly gain evidence for p until you eventually find yourself in a situation where belief is the only attitude that fits your total evidence? At some time t^1, your total evidence provides more support for suspending on p than it provides for any alternative doxastic attitude toward p. At some later time t^2, your total evidence provides more support for believing p than it provides for any alternative doxastic attitude toward p. What happens in between? As you gain evidence from t^1 to t^2, your evidence provides less and less support to suspension of judgment, and more and more support to belief. Is it plausible that, no matter how we fill in the details of this case, there can't be any point where belief and suspension *both* fit your total evidence?

Consider a true/false test where you get something valuable for each correct answer, lose something valuable for each incorrect answer, and neither gain nor lose anything at all for each question that you skip. Your total evidence might rule out checking the box marked 'F' while leaving you torn between skipping the question and checking the box marked 'T.' But intuitively, something similar could happen in a case where your evidence starts out providing more support for suspension than any alternative attitude and goes by gradual improvements to providing more support for belief than any alternative attitude. Somewhere along the way, your total evidence should provide equal support for belief and suspension while providing more support for each of these than for disbelief. Or, at least, it's hard to see why this couldn't happen. But permissive situations are possible if, anywhere along the way, your total evidence provides equal support for belief and suspension while providing more support for each of these than for disbelief. So it seems that permissive situations should be possible. And since DIE entails that permissive situations *aren't* possible, DIE seems false. It might, of course, turn out that that permissive situations actually *aren't* possible. But if nothing else, the *prima facie* plausibility of permissive situations gives anyone who would deny the possibility of permissive situations the burden of proof. These people owe *us* an argument for the conclusion that, even in a case where you go by gradual improvements in your evidence from the worst-case scenario to the best-case scenario, you'll never find yourself in a permissive situation.

Are there good arguments for this conclusion, then? Over the remainder of this paper, I will show why the answer is 'no.'

4. Uniqueness and Permissive Situations

DIE resembles some versions of the uniqueness thesis and there are powerful arguments for many versions of the uniqueness thesis. Since DIE entails that permissive situations aren't possible, someone might be tempted to argue against permissive situations by simply repeating some argument for uniqueness. In fact, however, no argument for any version of uniqueness causes any trouble for permissive situations.

In my view, the best arguments for one or another version of uniqueness come from Matheson (2011), White (2014), Dogramaci and Horowitz (2016), Greco and Hedden (2016), and Schultheis

(2017). But as I'll now argue, these arguments don't support any version of uniqueness that is both plausible and useful for rejecting permissive situations. Matheson's argument supports a version of uniqueness that clearly entails that permissive situations aren't possible, but the main premise in his argument has unacceptable consequences. His argument is simply this:

> For any proposition p and body of total evidence E, either E provides more support for p than for $\neg p$, it provides less support for p than for $\neg p$, or it supports p and $\neg p$ equally well. Belief uniquely fits E if E provides more support for p than for $\neg p$, disbelief uniquely fits E if E provides less support for p than for $\neg p$, and suspension uniquely fits E if E supports p and $\neg p$ equally well. So, for any proposition p and body of total evidence E, some doxastic attitude toward p uniquely fits E. (p. 365)

The second premise of this argument entails that suspending on p fits E only in the rare case where $\Pr(p|E) = 0.5$, but this conclusion is very hard to believe. Simply imagine a lottery with just three tickets, assume that I know this lottery will have exactly one winner, assume that I know that each ticket is equally likely to win, and assume that I don't have any other evidence relevant to the question whether a given ticket will win. Then, where 'E' names my total evidence, 't^1' names the proposition that ticket 1 will win, 't^2' names the proposition that ticket 2 will win, and 't^3' names the proposition that ticket 3 will win, my evidential situation looks like this.

$$\Pr(t^1 \vee t^2 \vee t^3 \,|\, E) = 1$$

$$\Pr(t^1 \,|\, E) = \Pr(t^2 \,|\, E) = \Pr(t^3 \,|\, E) = {}^1/_3$$

$$\Pr(\neg t^1 \,|\, E) = \Pr(\neg t^2 \,|\, E) = \Pr(\neg t^3 \,|\, E) = {}^2/_3$$

So in this case, my total evidence provides more support for each of $\neg t^1$, $\neg t^2$, and $\neg t^3$ than it provides for each of t^1, t^2, and t^3, and yet it *also* provides more support for $(t^1 \vee t^2 \vee t^3)$ than it provides for $\neg(t^1 \vee t^2 \vee t^3)$. So the major premise in Matheson's argument entails that believing that ticket 1 won't win uniquely fits my total evidence, believing that ticket 2 won't win uniquely fits my total evidence, believing that ticket 3 won't win uniquely fits my total evidence, and believing that either ticket 1, ticket 2, or ticket 3 *will* win also uniquely fits my total evidence. So, on Matheson's view, if you ask me what's going to happen and I say "ticket 1 won't win, ticket 2 won't win, and ticket 3 won't win, but either ticket 1, ticket 2, or ticket 3 *will* win," I can defend my answer by pointing out (a) that believing that ticket 1 won't win uniquely fits my total evidence, (b) that believing that ticket 2 won't win uniquely fits my total evidence, (c) that believing that ticket 3 won't win uniquely fits my total evidence, and (d) that believing that either ticket 1, ticket 2, or ticket 3 *will* win also uniquely fits my total evidence.

But this conclusion is beyond the pale. It's much worse than the conclusion that contradictory beliefs each uniquely fit your total evidence in standard lottery situations, since, in these situations, your credence in the proposition that a given ticket won't win should be *much* higher than my credence in each of $\neg t^1$, $\neg t^2$, and $\neg t^3$ (cf. Christensen 2004). Surely, rather than conclude that believing that ticket 1 won't win uniquely fits my total evidence, that believing that ticket 2 won't

win uniquely fits my total evidence, that believing that ticket 3 won't win uniquely fits my total evidence, and that believing that either ticket 1, ticket 2, or ticket 3 *will* win also uniquely fits my total evidence, it would be better to deny what many epistemologists already deny anyway: that suspending on p fits E only when $\Pr(p|E) = 0.5$. Once we deny this, however, there is no remaining idea in Matheson's argument that supports DIE.

What about the arguments for uniqueness from White, Dogramaci, Horowitz, Greco, Hedden, and Schultheis, then? I'm not entirely sure which view(s) these authors are defending. As Kopec and Titelbaum point out (2016: 190-92), at least sixteen different views fall under the label 'the uniqueness thesis,' and contributors to the debate often leave it implicit which view they have in mind. So far as I can tell, White, Dogramaci, Horowitz, Greco, Hedden, and Schultheis are all defending *credal* versions of uniqueness. If so, these versions of uniqueness might entail CIE (which I am already assuming is true), but they needn't conflict with the permissivist picture sketched in §3, since CIE needn't entail DIE and it might even *conflict* with DIE. At points, White seems interested in defending a doxastic version of uniqueness, but the doxastic version of uniqueness that he might be defending is just the negation of the following thesis. (I quote *verbatim*.)

> **Strong permissivism**: There are cases in which it is rationally permissible to believe p, but it is also rationally permissible to believe $\neg p$ instead, given the very same evidence. (2014: 312)

Obviously, the negation of this thesis doesn't entail DIE. In any case, White, Greco, and Hedden's arguments for their preferred versions of uniqueness at least *suggest* arguments for DIE, so it's worth discussing the relevant parts of their arguments.

According to what I will call White's 'toggling argument' (pp. 318-22), we should reject strong permissivism because it entails that a body of total evidence might put you in a situation where you can rationally flip-flop back and forth between believing p and believing $\neg p$ without any change in your evidence. Someone who found this argument persuasive might think that a similar argument shows that permissive situations aren't possible. The idea would be that, if belief in p and suspension of judgment with respect to p both fit your total evidence, then you can rationally flip-flop between believing p and suspending on p without any change in your evidence. Since this toggling could never be rational (the argument would continue), permissive situations must be impossible.

There are subtle issues here, but White's toggling argument doesn't threaten the possibility of permissive situations. The reason is simply that, on the stipulative definition of 'fit' that we have been working with throughout, there's no reason at all to think you can rationally flip-flop between believing p and suspending on p, without any change in your evidence, if belief in p and suspension of judgment with respect to p both fit your total evidence. Consider the true/false test from §3. Your total evidence might rule out checking the box marked 'F' while leaving you torn between skipping the question and checking the box marked 'T,' but it doesn't follow that you can rationally check the box marked 'T' and then erase your answer and then re-check the box marked 'T' and then re-erase your answer, round and round, without any change in your evidence. But exactly the

same point applies to the idea that, if permissive situations are possible, then you can rationally toggle between belief and suspension.[7]

What about the argument due to Greco and Hedden, then? Greco and Hedden start by defending a broadly Craigean (1990) account of our concept of rationality which says that, if you judge that someone's attitude toward p is rational, you thereby commit yourself to adopting that attitude toward p yourself, unless you think that you have relevant evidence that the person lacks (p. 372-3). Greco and Hedden then say that, *because* judging that someone's attitude toward p is rational involves this commitment, the following principle is true:

> **Deference**: Agents should satisfy the following conditional: If A judges both that S's attitude toward p is rational and that A doesn't have any relevant evidence that S lacks, then A adopts S's attitude toward p. (ibid)

Greco and Hedden don't say exactly *why* we should satisfy the conditional in Deference, but, presumably, they think we should satisfy this conditional because our attitudes will conflict with our commitments if we don't satisfy it. How do Greco and Hedden use Deference to support their version of uniqueness, then? Their argument is simply that it would be hard to explain why Deference is true if their version of uniqueness were false.

I'm not convinced that Deference *is* true, but I'll assume both that it is and that Greco and Hedden's argument for uniqueness works, since I want to show that the relevant analogue of their argument fails *even if* their argument works. According to the relevant analogue of Greco and Hedden's argument, it would be hard to explain why DJEF (below) is true if permissive situations were possible.

> **Deference for Judgments about Evidential Fit (DJEF)**: Agents should satisfy the following conditional: If A judges both that S's attitude toward p fits S's total evidence and that A doesn't have any relevant evidence that S lacks, then A adopts S's attitude toward p.

The problem with this argument is obvious—namely, people who aren't already convinced that permissive situations aren't possible will see no reason to accept DJEF, so DJEF can't support the conclusion that permissive situations aren't possible.

Here it's important to remember that, unlike the concept of rationality at work in Deference, the concept of evidential fittingness at work in DJEF is purely stipulative. As we defined the word 'fit' in §1, attitude A toward p fits body of total evidence E just in case the support that E provides to A is greater than *or equal to* the level of support that E provides to each of the alternative doxastic attitudes toward p. Because this is how we defined the word 'fit,' however, it's hard to see why judging both that someone's attitude toward p fits her total evidence and that you don't have

---

[7] Similar considerations apply to White's pill-popping case (315-17). Suppose you have two pills: one that will cause you to believe p and one that will cause you to suspend on p. If believing p and suspending on p both fit your total evidence, then, if you take one of the pills (and this doesn't change your relevant evidence), your resulting attitude will still fit your total evidence, on the relevant notion of 'fit.' Your resulting attitude might have *other* defects. (The pill might work by causing you to engage in wishful thinking without your *knowing* that you're engaged in wishful thinking, for example.) But one of its defects *won't* be that it doesn't fit your total evidence.

any relevant evidence that she lacks would require any commitment to adopting that attitude toward *p* yourself. After all, on this definition of 'fit,' judging that someone's attitude toward *p* fits her total evidence is consistent with your simultaneously judging that some *other* attitude toward *p* *also* fits her total evidence. So perhaps, given the way we defined the word 'fit,' judging both that someone's attitude toward *p* fits her total evidence and that you don't have any relevant evidence that she lacks would involve a commitment to adopting *one of the attitudes* that fit her total evidence. But of course, if multiple attitudes fit her total evidence, you can adopt one of the attitudes that fit her total evidence without adopting *her* specific attitude. So, even if judging both that someone's attitude toward *p* fits her total evidence and that you don't have any relevant evidence that she lacks would involve a commitment to adopting one of the attitudes that fit her total evidence, this by itself doesn't support DJEF.

This means that the general idea behind Greco and Hedden's argument for uniqueness doesn't support DIE. Since Matheson's argument relies on a false premise, and since Dogramaci, Horowitz, and Schultheis arguments only support *credal* versions of uniqueness, the best arguments for uniqueness in the literature cast no doubt on the possibility of permissive situations.

5. Akrasia and Permissive Situations

With this in mind, let's return to the primary concern of this paper. We started with the debate between doxastic voluntarists and doxastic involuntarists over the disanalogy thesis—the thesis that beliefs differ from actions in at least the following way: while cognitively healthy people often exhibit direct control over their actions, there is no possible scenario where a cognitively healthy person exhibits direct control over her beliefs. As we saw, Raz, Ginet, Frankish, Nickel, and McHugh all reject the disanalogy thesis, and they all reject it on the basis of permissive situations. According to Kurt Sylvan (2016), however, these authors are all mistaken, since permissive situations aren't possible. Sylvan starts his argument against permissive situations with a principle that he calls the "Underconfidence Datum." I quote *verbatim*:

> **(Underconfidence Datum)** It is less than fully rational to believe at *t* that the evidence for *p* is sufficient but to remain agnostic on *p* at *t*. This amounts to akratic underconfidence. (p. 1648)

Sylvan then says that, because the Underconfidence Datum is true, the following wide-scope principle is also true (again, I quote *verbatim*):

> **(REQ)** Rationality requires that: if one believes the evidence for *p* is sufficient at *t*, one is not agnostic about *p* at *t*. (ibid)

From here Sylvan argues that, if REQ is true, then it's also true that

> **(P)** [i]f one's evidence for *p* is sufficient at *t*, one lacks sufficient epistemic reason to be agnostic on *p* at *t*. (p. 1650)

Finally, Sylvan argues that, if P is true, then DSIC below is *false*.

(DSIC) When a person's evidence for *p* at *t* is sufficient but not conclusive, it would be epistemically rational for this person to believe *p* at *t* but also epistemically rational for her to be agnostic on *p* at *t* instead. (p. 1648)

According to Sylvan, the negation of DSIC entails that permissive situations aren't possible, since cases of sufficient but inconclusive evidence would be the only cases where permissive situations *could* be possible (p. 1650). I find a lot in Sylvan's argument compelling. But as Sylvan makes clear (ibid), his argument depends on the Underconfidence Datum, and as I'll now argue, the Underconfidence Datum can't support the conclusion that permissive situations aren't possible.

Consider the difference between what we might call 'balanced cases,' where you think the evidence for *p* is sufficient because you think it provides *equal* support to belief and suspension, and what we might call 'belief-favoring cases,' where you think the evidence for *p* is sufficient because you think it provides *more* support for belief than suspension. Is the Underconfidence Datum supposed to apply just to belief-favoring cases, or is it supposed to apply to both belief-favoring cases and balanced cases? If the latter, then the Underconfidence Datum seems no *datum* at all. Instead, it seems clearly false. After all, since there needn't be anything incoherent about suspending on *p* while thinking that belief in *p* and suspension with respect to *p* both fit your total evidence, since this combination of attitudes needn't manifest any weakness of the will, since you needn't be acting against your better judgment if you suspend on *p* while thinking that belief in *p* and suspension with respect to *p* both fit your total evidence, and so on, there needn't be anything akratic about suspending on *p* while thinking that belief in *p* and suspension with respect to *p* both fit your total evidence.

If the Underconfidence Datum is supposed to apply just to belief-favoring cases, however, then we get an invalid argument unless we make the requisite changes to REQ and P, and subsequently modify DSIC as follows:

(DSIC*) When at *t* a person's evidence for *p* provides more support for believing *p* than for agnosticism with respect to *p*, it would be epistemically rational for this person to believe *p* at *t* but also epistemically rational for her to be agnostic on *p* at *t* instead.

Of course, DSIC* *does* seem false. But the negation of DSIC* clearly doesn't that permissive situations aren't possible, since the negation of DSIC* just entails that, in at least some situations that *aren't* permissive situations, it's false that it would be epistemically rational for the relevant person to believe *p* but also epistemically rational for this person to suspend on *p*. So either the Underconfidence Datum is false, or, on the correct way to understand the Underconfidence Datum, it doesn't entail that permissive situations aren't possible. Either way, permissive situations survive Sylvan's worries about akrasia.[8]

---

[8] It's worth noting that, the way Sylvan explicitly characterizes agnosticism as a settled resistance to belief that isn't consistent with an intention to give further consideration to the relevant evidence (pp. 1652-3), agnosticism is clearly not *equivalent* to suspension of judgment. Instead, it seems clearly a *species* of suspension of judgment, in roughly the way that dogmatism (defined as belief conjoined with an intention to avoid further consideration of the relevant evidence) is a species of belief. But of course, on *this* characterization of agnosticism, the truth-value of DSIC is clearly irrelevant to the possibility of permissive situations, as I characterized them in §3.

6. Vagueness and Permissive Situations

In response to the view sketched in §3, several readers have suggested that, because we get a sorites sequence as your evidential situations improves, vagueness threatens and we can't draw any of the conclusions we might be tempted to draw. Is vagueness a problem for the view sketched in §3, then? Perhaps. But at most, it's only a superficial problem, since vagueness only *supports* Raz, Ginet, Frankish, Nickel, and McHugh's arguments against the disanalogy thesis.

According to the view sketched in §3, in any case where your evidence starts out providing more support for suspension than belief and goes by sufficiently gradual improvements to providing more support for belief than suspension, there should be some point where suspension and belief both fit your total evidence. With this in mind, consider the following thesis.

> **One Penumbral Zone (OPZ)**: As your evidence improves, a single penumbral zone will separate the situations where suspension uniquely fits your total evidence from the situations where belief uniquely fits your total evidence. In graphic form:

Suspension uniquely
fits your total
evidence
                                                                    Belief uniquely fits
                                                                    your total evidence
———
                              Penumbral
                              zone
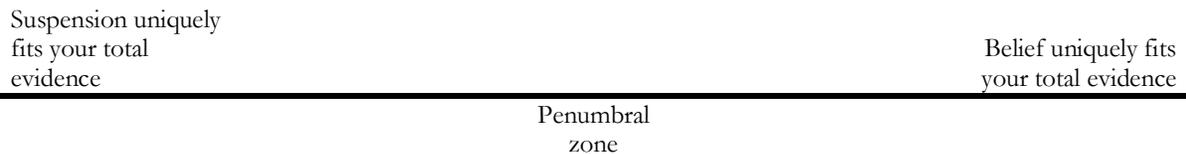
On a supervaluationist interpretation of OPZ, this thesis says that you will go from a range of situations where the sentence 'suspension uniquely fits your total evidence' is true under every admissible precisification to a range of situations where this sentence is true under some admissible precisifications and false under others, and you will eventually find yourself in a range of situations where this sentence is false under every admissible precisification. According to fuzzy semantics, OPZ says you will go from a range of situations where the sentence 'suspension uniquely fits your total evidence' is completely true to a range of situations where it's partially true and partially false, and you will eventually land in a range of situations where it's completely false. On an epistemicist interpretation, OPZ says you will go seamlessly from a range of situations where the sentence 'suspension uniquely fits your total evidence' is true to a range of situations where it's false, but in the center of the penumbral zone, you will be hopelessly unable to tell which of these situations you are in. And so on. Thus, if OPZ is true, some of the most prominent theories of vagueness entail that, as your evidential situation improves, you will *never* find yourself in a situation where it's true that belief and suspension both fit your total evidence.

Does OPZ threaten the view sketched in §3, then? Not really. First, whatever theory of vagueness turns out correct, it's hard to see why anyone should accept OPZ instead of TPZ, below.

> **Two Penumbral Zones (TPZ)**: As your evidence improves, one penumbral zone will separate the situations where suspension uniquely fits your total evidence from the situations where suspension and belief *both* fit your total evidence, and a second penumbral

zone will separate the situations where suspension and belief both fit your total evidence from the situations where belief uniquely fits your total evidence. In graphic form:

Suspension uniquely fits your total evidence	Suspension and belief both fit your total evidence	Belief uniquely fits your total evidence
Penumbral zone		Penumbral zone
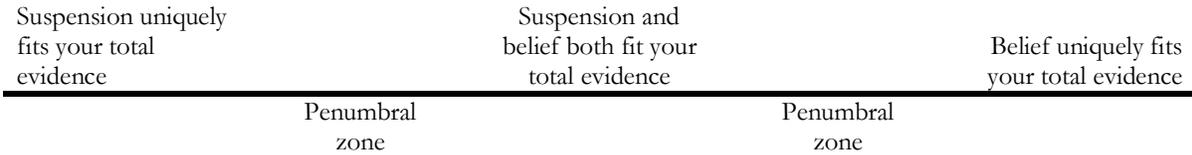
OPZ and TPZ are *both* consistent with every theory of vagueness—there are supervaluationist versions of TPZ, fuzzy semanticist versions of TPZ, epistemicist versions of TPZ, and so on—so the arguments for and against the various theories of vagueness are neutral between OPZ and TPZ. But now two things need emphasis. First, whatever theory of vagueness turns out correct, TPZ seems more plausible than OPZ, since TPZ seems relevantly analogous to the completely uncontroversial claim that, as we move around the perimeter of the color wheel from (say) blue to yellow, we pass through a green region, and we get penumbral zones *before* and *after* the green region—between the blue and green regions, and then again between the green and yellow regions. After all, as we stipulatively defined the word 'fit' in §1, suspension and belief both fit your total evidence just in case a cognitively healthy human being who was considering the same question and who had exactly your total evidence would experience an attraction toward suspension that would feel exactly as strong as her attraction toward belief. The claim that *this* couldn't happen seems about as plausible as the claim that there's no such color as green.

Second—and more important—even if OPZ is true, it conflicts with the letter but not the spirit of the view sketched in §3. Call the zone where TPZ says that suspension and belief both fit your total evidence 'the disputed zone,' let 'SU' name the sentence 'suspension uniquely fits your total evidence,' and let 'BU' name the sentence 'belief uniquely fits your total evidence.' Then, if OPZ is true and you are in the center of the disputed zone, supervaluationists will have no principled reason to reject the suggestion that SU and BU are true under an equal number of admissible precisifications; fuzzy semanticists will have no principled reason to reject the suggestion that SU and BU are true to exactly the same degree; epistemicists will be forced to say that exactly one of SU and BU is true but we're hopelessly unable to tell which one it is; and so on. Even on Cian Dorr's (2003) view, which entails that it's neither determinately true that you don't know that SU is true nor determinately true that you don't know that BU is true, it's still neither determinately true that you *do* know that SU is true nor determinately true that you *do* know that BU is true (even if you're silly enough to *believe* either that SU is true or that BU is true).[9] In the middle of the disputed zone, even proponents of OPZ must admit some deep (possibly hard to describe) symmetry between SU and BU.

---

[9] Of course Dorr's view only entails this conclusion *if* OPZ is true. For if TPZ is true, even *his* view is consistent with the conclusion that it *is* determinately true both that you don't know that SU is true and that you don't know that BU is true. Given that it seems patently obvious that you can't know either that SU is true or that BU is true if you're in the center of the disputed zone, I'm guessing Dorr would plump for TPZ. (Thanks to Daniel Nolan and Brian Cutter for bringing this article to my attention.)

Why is this important? Because it highlights the fact that, for purposes of evaluating Raz, Ginet, Frankish, Nickel, and McHugh's respective arguments against the disanalogy thesis, it doesn't matter whether OPZ is true. I won't burden readers with the details of their arguments, since, for our purposes, we can simply note something that Nickel (for example) makes clear: permissive situations matter for evaluating the disanalogy thesis because they give us examples of good evidence that isn't fully compelling (2016: 315-18). But of course, in the center of the disputed zone, we get exactly the *same* examples of good evidence that isn't fully compelling. OPZ and TPZ give different answers to the question whether you're in a *permissive situation* if you're in the center of the disputed zone, but they don't give different answers to the question that matters for evaluating Raz, Ginet, Frankish, Nickel, and McHugh's arguments against the disanalogy thesis: whether you have *good evidence that isn't fully compelling* if you're in the center of the disputed zone. OPZ and TPZ both *agree* that it's possible to be in the center of the disputed zone, and they both *agree* that, if you *are* in the center of the disputed zone, then you will have good evidence that isn't fully compelling. This means that, anytime you find yourself in the middle of the disputed zone, you will find yourself in a situation that's permissive *in at least the sense that matters for evaluating the disanalogy thesis*. Thus, even if OPZ turns out true—and again, OPZ strikes me as about as plausible as the suggestion that there's no such color as green—vagueness won't threaten the kinds of permissive situations that matter for evaluating the disanalogy thesis.

7. Conclusion

What should we think about these kinds of permissive situations, then? They seem obviously, uncontroversially possible. Of course, it doesn't follow that the disanalogy thesis is false. But it does follow that doxastic involuntarists can't respond to the arguments from Raz, Ginet, Frankish, Nickel, McHugh, and others by arguing that permissive situations aren't possible. To resist these arguments, doxastic involuntarists must show that a cognitively healthy person can't exercise direct control over whether she believes *p even in a situation where she has good evidence for p that isn't fully compelling*. Perhaps doxastic involuntarists can do this, but they haven't done this yet. And this fact is remarkable. Given that doxastic involuntarism is virtually orthodoxy among contemporary epistemologists—as Steup (2000), Booth (2007), Hieronymi (2009), Boyle (2009), Peels (2015), and many others point out—it's hard to believe that formidable arguments for doxastic *voluntarism* haven't been answered.[10]

References

Bennett, Jonathan (1990), "Why is Belief Involuntary?," *Analysis*, Vol. 50, No. 2: pp. 87-107.

Bergman, Michael (2006), *Justification without Awareness* (Oxford: Oxford University Press).

Berker, Selim (2013), "Epistemic Teleology and the Seperateness of Propositions," *The Philosophical Review*, Vol. 122, No. 3: pp. 337-93.

Booth, Anthony Robert (2007), "Doxastic Voluntarism and Self-Deception," *Disputatio*, Vol. 2, No. 22: pp. 115-30.

Boyle, Matthew (2009), "Active Belief," *Canadian Journal of Philosophy*, Vol. 35, No. 1: pp. 119-47.

Buchak, Lara (2014), "Belief, Credence, and Norms," *Philosophical Studies*, Vol. 116: pp. 285-311.

Christensen, David (2004), *Putting Logic in its Place* (Oxford: Oxford University Press).

Clarke, Roger (2013), "Belief is Credence One (In Context)," *Philosophers' Imprint*, Vol. 13, No. 11: pp. 1-18.

Cohen, L. Jonathan (2010), "Rationality," in *A Companion to Epistemology*, ed. J. Dancy, E. Sosa, and M. Steup (Oxford: Wiley-Blackwell).

Cohen, Stewart (2016), "Reasons to Believe and Reasons to Act," *Episteme*, Vol. 13, No. 4: pp. 427-38.

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

Dogramaci, Sinan and Sophie Horowitz (2016), "An Argument for Impermissivism about Evidential Support," *Philosophical Issues*, Vol. 26: pp. 130-47.

Dorr, Cian (2003), "Vagueness without Ignorance," *Philosophical Perspectives*, Vol. 17: pp. 83-113.

Fantl, Jeremy and Matthew McGrath (2009), *Knowledge in an Uncertain World* (Oxford: Oxford University Press).

Feldman, Richard (2001), "Voluntary Belief and Epistemic Evaluation," in *Knowledge, Truth, and Duty*, M. Steup ed. (Oxford: Oxford University Press).

Frankish, Keith (2007), "Deciding to Believe Again," *Mind*, Vol. 116, No. 2: pp. 523-47.

Friedman, Jane (2013), "Rational Agnosticism and Degrees of Belief," *Oxford Studies in Epistemology*, Vol. 4: pp. 57-81.

Ganson, Dorit (2008), "Evidentialism and Pragmatic Constraints on Outright Belief," *Philosophical Studies*, Vol. 139: pp. 441-58.

Ginet, Carl (2001), "Deciding to Believe," in *Knowledge, Truth, and Duty*, M. Steup ed. (Oxford: Oxford University Press).

Greco, Daniel and Brian Hedden (2016), "Impermissivism and Metaepistemology," *The Journal of Philosophy*, Vol. 103, No. 8: pp. 365-95.

Hieronymi, Pamela (2009), "Believing at Will," *Canadian Journal of Philosophy*, Vol. 35: pp. 149-87.

Kopec, Matthew and Michael Titelbaum (2016), "The Uniqueness Thesis," *Philosophy Compass*, Vol. 11, No. 4: pp. 189-200.

Leitgeb, Hannes (2014), "The Stability Theory of Belief," *The Philosophical Review*, Vol. 123, No. 2: pp. 131-71.

Locke, Dustin (2014), "The Decision-Theoretic Lockean Thesis," *Inquiry*, Vol. 57, No. 1: pp. 28-54.

Matheson, Jonathan (2011), "The Case for Rational Uniqueness," *Logos and Episteme*, Vol. 2: pp. 359-73.

McHugh, Conor (2014), "Exercising Doxastic Freedom," *Philosophy and Phenomenological Research*, Vol. 88, No. 1: pp. 1-37.

Nickel, Philip (2010), "Voluntary Belief on a Reasonable Basis," *Philosophy and Phenomenological Research*, Vol. 81, No. 2: pp. 312-34.

Peels, Rik (2015), "Believing at Will is Possible," *Australasian Journal of Philosophy*, Vol. 93, No. 3: pp. 524-41.

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

Raz, Joseph (1999), *Engaging Reasons* (Oxford: Oxford University Press).

Rinard, Sussana (2017), "No Exception for Belief," *Philosophy and Phenomenological Research*, Vol. 94, No. 1: pp. 121-43.

Ross, Jacob and Mark Schroeder (2014), "Belief, Credence, and Pragmatic Encroachment," *Philosophy and Phenomenological Research*, Vol. 88, No. 2: pp. 259-88.

Rubin, Katherine (2015), "Total Pragmatic Encroachment and Epistemic Permissiveness," *Pacific Philosophical Quarterly*, Vol. 96: pp. 12-38.

Schultheis, Ginger (2017), "Living on the Edge: Against Epistemic Permissivism," *Mind*, DOI: 10.1093/mind/fzw065.

Sosa, Ernest (2010), "Value Matters in Epistemology," *The Journal of Philosophy*, Vol. 107, No. 4: pp. 167-90.

Staffel, Julia (2016), "Beliefs, Busses, and Lotteries," *Philosophical Studies*, Vol. 173: pp. 1721-34.

Steup, Matthias (2000), "Doxastic Voluntarism and Epistemic Deontology," *Acta Analytica*, Vol. 15, No. 1: pp. 25-56.

——— (2017), "Believing Intentionally," *Synthese*, DOI: 10.1007/s11229-015-0780-7.

Sylvan, Kurt (2016), "The Illusion of Discretion," *Synthese*, Vol. 193: pp. 1,635-1,665.

Weatherson, Brian (2012), "Knowledge, Bets and Interests," in *New Essays on Knowledge Ascriptions*, eds. J. Brown and M. Gerken (Oxford: Oxford University Press).

White, Roger (2014), "Evidence Cannot be Permissive," in *Contemporary Debates in Epistemology*, M. Steup, J. Turri, and E. Sosa eds. (Oxford: Wiley Blackwell).

Worsnip, Alex (2015), 'The Conflict of Evidence and Coherence', *Philosophy and Phenomenological Research*, DOI: 10.1111/phpr.12246.