



Arbitrary switching and concern for truth

Sophia Dandelet¹

Received: 2 January 2023 / Accepted: 31 August 2023

© The Author(s) 2023

Abstract

This essay is about a special kind of transformative choice that plays a key role in debates about permissivism, the view that some bodies of evidence permit more than one rational response. A prominent objection to this view contends that its defender cannot vindicate our aversion to arbitrarily switching between belief states in the absence of any new evidence. A prominent response to that objection tries to provide the desired vindication by appealing to the idea that arbitrary switching would involve a special kind of transformative choice: the choice to change one’s epistemic standards, i.e., one’s commitments regarding the relative importance of achieving true belief and avoiding false belief. My first aims here are to argue that this response is unsuccessful and propose an alternative. My secondary aim is to consider how this discussion might bear on more general debates about transformative choice.

Keywords Permissivism · Arbitrariness · Transformative experience · Epistemic utility theory · Epistemic consequentialism · One thought too many

1 Introduction

A *personally transformative experience* is an experience that “changes you enough to substantially change your point of view, thus substantially revising your core preferences or revising how you experience being yourself” ((Paul, 2014) p. 16). Sometimes, such experiences take you by surprise. But in other cases, you find yourself facing a decision about whether to undergo an experience that you recognize as one that would

Many thanks to Jane Friedman, Alex Kerr, Ram Neta, Richard Pettigrew, R. Jay Wallace, Tim Williamson, Snow Zhang, two anonymous reviewers at *Synthese*, and audience members at the Workshop on Transformative Experience at Yale in April 2022—for comments, questions, and conversations that greatly improved the manuscript.

✉ Sophia Dandelet
scd58@cam.ac.uk

¹ University of Cambridge, Cambridge CB2 3HU, UK

change your core commitments. Let's say that when you knowingly decide to undergo a personally transformative experience, you make a *transformative choice*.

This essay is about a special kind of transformative choice that plays a key role in debates about *epistemic permissivism*. Very roughly, epistemic permissivism is the view that some bodies of evidence permit more than one rational response. A prominent objection to this view contends that its defender cannot locate a distinctively epistemic reason, accessible from the first-person point of view, to avoid arbitrarily switching between belief states in the absence of any new evidence. A prominent response to that objection—call it the *standards-based response*—tries to make the existence of such a reason consistent with permissivism by appealing to the idea that arbitrary switching would involve a special kind of transformative choice: the choice to change one's *epistemic standards*, i.e., one's commitments regarding the relative importance of achieving true belief and avoiding false belief.

My first aims here are to argue that the standards-based response is unsuccessful and propose an alternative. In a nutshell, the problem that I see with the standards-based response is that it raises questions about the nature of epistemic standards that have no satisfying answers. The alternative that I put forward preserves a core idea of the standards-based response—namely, that our interest in avoiding arbitrary switching stems from a concern for truth—but rejects the consequentialist analysis of concern for truth that the standards-based response implicitly endorses, according to which a person's concern for truth is constituted by her way of valuing the twin epistemic goals of achieving true belief and avoiding false belief. On the analysis of concern for truth that I prefer, having a concern for truth involves occupying a blinkered deliberative perspective in which you are insensitive to the considerations bearing on the prudential value of your belief states, and from which you approve of this insensitivity. As we'll see, this analysis provides a distinctively epistemic vindication of our aversion to arbitrary switching that is consistent with permissivism, and without positing the existence of epistemic standards.

My secondary aim is to consider how these discussions might bear on more general debates about transformative choice. In giving my preferred account of what it is to manifest concern for truth, I draw on Bernard Williams' insight that some of our most central commitments involve deliberative insensitivity to certain reasons. If this is right, it has a number of implications for debates about transformative choice. First, it illustrates that our commitments might be structurally very diverse, and that therefore it may not make sense to demand a unified theory of how to make potentially transformative decisions. Relatedly, because it implies that some of our commitments cannot be fully understood as commitments to ends, it reveals a limitation of investigating transformative choice in terms of means-end (i.e., prudential) rationality.

2 Permissivism and the arbitrariness objection

Elizabeth and Jane have examined the same evidence and come to different conclusions. Elizabeth thinks that Mr. Darcy has wronged Mr. Wickham deliberately and maliciously, while Jane thinks that there isn't yet enough evidence to establish this

distressing conclusion.¹ According to the *permissivist*, it might be that Jane and Elizabeth both believe permissibly. The *uniquer* disagrees. She thinks that there is just one permitted belief state for each body of evidence, so that if one of Jane and Elizabeth believes permissibly, the other must believe impermissibly.

2.1 Some apparent advantages of permissivism

At least at first glance, permissivism seems to have several advantages. It straightforwardly respects the commonsense idea that reasonable people can disagree even when presented with the same evidence ((Rosen, 2001) p. 71). Relatedly, it provides an attractive “out” for someone who doesn’t share the belief of a respected interlocutor, yet finds that persistent attempts to defend her own position seem to beg the question or hit justificatory bedrock prematurely.² If this person is a permissivist, she can say that her position and her interlocutor’s are both permitted and leave it at that. She doesn’t need to take the dogmatic stance of insisting without reason that her own position is better.³

Permissivism is also attractive because it preserves a certain symmetry between epistemic rationality, on the one hand, and prudential rationality and morality, on the other. Any plausible theory of prudence or ethics will allow that in at least some circumstances, the prudential/ethical considerations allow more than one course of action. For instance, if you are in the position of Buridan’s ass, the reasons of prudence permit you to go to either bushel of hay. If you have promised someone cake without specifying the flavor, then (other things being equal) the reasons of morality permit you to give her chocolate cake and also permit you to give her yellow cake. Permissivism says that epistemic rationality is analogous to prudence and morality in this respect: in at least some circumstances, the epistemic considerations permit more than one response.

2.2 The arbitrariness challenge

The problem with permissivism is that it seems to infect our beliefs with a kind of *arbitrariness*. There are a number of ways to develop this line of objection. On what I take to be the most straightforward of these, the permissivist is accused of not being able to explain why switching between belief states without any new evidence—what we can call *arbitrary switching*—is unacceptable from a purely epistemic point of view.⁴

¹ The example is adapted from Austen (1918).

² This is mentioned in Christensen (2007) as a possible advantage of permissivism, although Christensen rejects permissivism in the end for reasons that are similar to White’s in White (2005).

³ I should note that there might be ways for the naive non-permissivist to accommodate these considerations that seem at first glance to support naive permissivism. For suggestions along these lines, see Kelly (2014) and Greco and Hedden (2016). But I won’t say more about this, since my primary aim here is to consider the relative merits of naive permissivism and standards permissivism, rather than the relative merits of permissivism and non-permissivism.

⁴ The objection is first articulated, to my knowledge, in White (2005) pp. 448–9.

The problem can be brought out as follows: Suppose that you have carefully examined the evidence and thereby arrived at agnosticism about whether God exists. You are also persuaded by Pascal's arguments that the thing to do, all things considered, is to embark on a project of self-indoctrination to get yourself to believe in God. Nevertheless, you feel what Williams called "a strong internalized objection" to this idea ((Williams, 1970) p. 150). For although the contemplated project seems to you, from the practical perspective, like the very best course of action, it also seems that if you restrict your attention to your distinctively epistemic reasons—roughly, the reasons that stem from your concern for truth, knowledge, understanding, or some other epistemically-flavored good—these decisively weigh against believing in God *and* taking steps to make yourself believe. More generally, if you restrict your attention to the epistemic reasons, it seems to you that you ought not switch out of the belief states that you have arrived at by examining the evidence unless and until you get new relevant evidence. We would like to be able to explain why this is.

One natural way of doing this is to suppose that every body of evidence permits just one belief state and that people generally recognize this, even if they would not put the point in exactly those terms. If these things are true, then we can easily understand why people feel a strong, epistemically-flavored objection to projects like the one that Pascal recommends. The permissivist, however, cannot appeal to this explanation because it involves the denial of her view. So what are her options? She could (1) say that everyone mistakenly believes in uniqueness—but this will not be attractive to her if she wants to avoid an error theory. She could (2) say that, although permissivism is true, people can never tell when they are in a permissive state—but this would substantially diminish the interest of her view. Or, she could (3) give some distinctively epistemic vindication of our aversion to arbitrary switching that applies even in known permissive cases. But how would this vindicating explanation go, exactly? The arbitrariness challenge, as I understand it, demands that the permissivist answer this question if she wants to embrace option (3).

I should flag that my version of the arbitrariness challenge is much less forceful than the standard one, which aims to show that permissivism is either false, parasitic on an error theory, or uninteresting by assuming—as I do not—that option (3) is off the table.⁵ Let me quickly say why I do not like this assumption, in part to justify my somewhat idiosyncratic presentation of the objection, but also because these concerns serve to introduce a theme that I will continue to develop later.

2.3 Acknowledging the nebulousness of the challenge

What would entitle us to declare that option (3) is off the table? That is, what would it take to establish that, if you knew that arbitrary switching would leave you in a permissible belief state, you could have no reasonable and distinctively epistemic objection to it? I think that a satisfying defense of this claim would involve at least

⁵ "For suppose...I am persuaded that there is this range of rationally permissible degrees of conviction. If I am correct in really thinking this, then there should be nothing wrong with arbitrarily choosing a verdict...For there should be nothing wrong with my arbitrarily choosing a degree of conviction (included perhaps by a magic pill) that is within the rationally permissible range" ((White, 2005) p. 453).

two things: first, a well-developed analysis of epistemic normativity, one that helps us see what distinguishes it from other kinds of normativity; and second, an argument showing that the *only* vindication of our aversion to arbitrary switching that can be spelled out in terms of distinctively epistemic considerations is one that invokes the value of avoiding epistemically impermissible states.

These things are necessary because we simply don't have a clear and widely shared understanding about what makes a rationalizing explanation of some attitude or action distinctively epistemic. We can gesture at the target concept by saying that an epistemic rationalizing explanation is one that appeals to the subject's concern for truth, knowledge, understanding, accuracy, or some other epistemically-flavored good. But of course, this informal gloss will not help us make fine-grained judgments about when someone's actions or attitudes can be justified in a distinctively epistemic way—both because there are potentially many epistemically-flavored goods, and because there are potentially many different ways to manifest concern for them. In light of this, it seems to me that we will not be entitled to treat the arbitrariness worry as an argument against permissivism until we have a richer understanding of the epistemic normative domain.⁶

Better, then, to treat the arbitrariness worry as an interesting explanatory challenge for the view. Again, that challenge is to (1) give a vindicating, non-error theoretic explanation of why we might be averse to arbitrary switching even in known permissive cases, and (2) do so in terms that have a plausible claim to be called *epistemic*.

3 The standards-based response

So, can the permissivist meet the arbitrariness challenge? It is sometimes suggested that she can—as long as she abandons the original version of her view in favor of *standards permissivism*.

3.1 Standards permissivism

The basic idea of standards permissivism is that whether a belief state is rational for you depends not only on your evidence, but also on your preferred way of balancing the twin epistemic aims of believing truly and not believing falsely. In other words,

⁶ I am not the first to point out that the domain of *epistemic normativity* is a philosophical invention whose borders are contested. For helpful discussions of this point, see Cohen (1984) and Friedman (2020). One point of contention that is particularly relevant for our purposes is about whether there are epistemic reasons for action. When some philosophers talk about epistemic reasons, they are talking about reasons for belief rather than reasons for action—where these two categories are understood to be completely separate in an important way. (For developments of this position, see Hieronymi (2004) and Hieronymi (2005).) But other philosophers want to say that there is a normatively significant unity between our evidential reasons for belief and certain kinds of reasons for action (See Singer (forthcoming) for discussion.) This debate matters for present purposes because you might think that arbitrary switching is only possible through some kind of action (e.g., meditation or magic pill-popping), in which case you might also think that the only reasons that can weigh for or against it are reasons for action. If this is your position, then you will want to spell out the arbitrariness objection in terms of reasons for action, and the intelligibility of the objection will depend on there being distinctively epistemic reasons for action.

the standards permissivist treats epistemic permissibility not as a two-place relation between a body of evidence and a belief state, but as a three-place relation between a body of evidence, a belief state, and a way of balancing the twin epistemic aims. This third relatum is sometimes called an *epistemic standard*.

According to the standards permissivist, what is rational for you to believe depends on your epistemic standard in roughly the following way: the more weight your epistemic standard gives to not believing falsely, the more evidence for p is required in order for you to be rational in believing that p , and the more weight your epistemic standard gives to believing truly, the less evidence for p is required in order for you to be rational in believing p .⁷ For an illustration of this idea, think back to Jane and Elizabeth. Let's suppose that, while Jane and Elizabeth agree that believing truly and not believing falsely are both important to some degree, they disagree about the relative importance of these two goals. Jane thinks that avoiding error is extremely important. Elizabeth, by contrast, is more intellectually adventurous. She knows that avoiding error is important to some extent, but she puts a higher premium on true belief than Jane does. According to the standards permissivist, because Elizabeth and Jane value the twin epistemic goals in different ways, what is rational for them to believe might be different—even in cases where they have all the same evidence. Jane, being more intellectually cautious, will require more evidence for p in order to be rational in believing p ; Elizabeth, being more intellectually adventurous, will require somewhat less.

The standards permissivist needn't claim that any belief state-evidence pairing can be rationalized by some epistemic standard. It is open to her to insist that a *reasonable* epistemic standard—one such that believing in line with it yields epistemically permissible beliefs—will never license, e.g., believing that p when p is 51% likely on the evidence. (Although of course, if she says this, she will eventually need to justify her claim by appealing to a systematic account of the difference between reasonable and unreasonable standards, and she must also make sure that her account of this difference renders intelligible the normative importance that she attributes to it.) She does need to say, however, that there is sufficient diversity in the realm of reasonable epistemic standards to allow that two people might be rational in occupying different belief states on the same evidence. Otherwise, her view wouldn't count as a kind of permissivism.⁸

⁷ There are many ways to model epistemic standards and their relation to epistemic permissibility, and which you choose will depend on a number of factors. If you are trying to develop a version of standards permissivism for full belief, you may want to model epistemic standards using *accuracy scores* which are essentially specialized utility functions that assign scores to belief states based on how closely they approach the truth. This is the strategy pursued by Schoenfield in Schoenfield (2012) and Schoenfield (2022). However, Pettigrew and Horowitz have argued convincingly that if you are trying to develop a version of standards permissivism for credences rather than full belief states, you may want to model epistemic standards using some other formal object, e.g., prior probabilities or a rule for selecting prior probabilities (see Horowitz (2017) p. 276 and Pettigrew (2022) p. 32–44). For my purposes, it won't matter which of these frameworks we choose, because the problem that I will eventually raise for the standards-based solution is not a problem for any particular way of modeling epistemic standards. Rather, it is a problem that we face as we try to understand the nature of the attitude that a person is supposed to take towards her own epistemic standard.

⁸ Here is how Kelly sums up these ideas: “So long as there are at least some possible cases in which it is reasonable for different individuals to give at least somewhat different weights to the [twin epistemic] goals, then this can affect how much evidence they should hold out for before they take up the relevant belief.

Another tenet of standards permissivism is the idea that every reasonable epistemic standard is such that, for any body of evidence, there is exactly one belief state that is recommended by this standard. So, even if Jane and Elizabeth are permitted to occupy different belief states, neither is such that agnosticism and belief in Mr. Darcy's malice are both permissible for her.

In sum, then, the standards permissivist is someone who believes that:

*Standards permissivism*⁹

1. A belief state B is epistemically permitted for a person X iff: B is recommended by X 's epistemic standard in light of X 's evidence, and X 's epistemic standard is reasonable.
2. For some body of evidence E and incompatible belief states B and B' : one reasonable epistemic standard recommends B in light of E , while another reasonable epistemic standard recommends B' in light of E .
3. No reasonable epistemic standard recommends two incompatible belief states in light of the same body of evidence.

Thanks to claim (2), standards permissivism implies that perfectly permissible disagreement is possible, even amongst interlocutors with all the same evidence. It has this much in common with original permissivism. But standards permissivism diverges from its ancestor because of claim (3), which together with (1) entails that for any person X , body of evidence E , and incompatible doxastic states: at most one of these two states is permitted for X in light of E . This position is sometimes called *intrapersonal uniqueness*.

3.2 Immodesty

You might think that worries about arbitrariness go away as soon as we accept intrapersonal uniqueness.¹⁰ But, as several people have pointed out, this is too quick.¹¹ For if we suppose that an arbitrary switching event could change your epistemic standard as well as your beliefs, then there are situations in which arbitrary switching would bring you to an epistemically permitted belief state. So, absent some reason for thinking that a person's standards are an unalterable fact about her, the standards permissivist

Footnote 8 continued

There will then be possible bodies of evidence that fall within the relevant margin, bodies of evidence relative to which belief is a perfectly reasonable response on the part of the person who is somewhat more concerned to believe the truth, and relative to which suspension of judgment is a perfectly reasonable response on the part of the person who is somewhat more concerned to avoid believing what is false" ((Kelly, 2014) p. 302). Note that Kelly himself has certain reservations about this picture, stemming from concerns about whether epistemic rationality is really best understood as a kind of instrumental rationality ((Kelly, 2014) endnote 3).

⁹ Something like this view is articulated in Kelly (2014) pp. 301–303, Pettigrew (2022) pp. 22–31, Schoenfield (2012), and Schoenfield (2022).

¹⁰ In Kelly (2014), Kelly interprets White's arbitrariness worry as a worry for intrapersonal permissivism that will only make trouble for interpersonal permissivism if the latter implies the former.

¹¹ See, for instance, ((White, 2005) p. 451) and ((Schoenfield, 2012) p. 12)

cannot vindicate an aversion to arbitrary switching in known permissive cases by suggesting that such switching will always seem, from the first-person perspective, to lead to impermissible believing.

How, then, is standards permissivism supposed to address the arbitrariness challenge? In brief, the standards permissivist vindicates our aversion to arbitrary switching by appealing to the hypothesis that the belief states recommended by your own epistemic standard will always seem to you better than the belief states recommended by any alternative standard—not because they seem to you like the only rational belief states that are possible for you, but because they seem to you to do the best job of balancing the twin epistemic goals of believing truly and not believing falsely. This is sometimes called the *immodesty* hypothesis. On this picture, the reason that you take yourself to have to avoid arbitrary switching has nothing to do, in fact, with epistemic rationality. It stems from your perception that your own epistemic standard corresponds to the best way of balancing the twin epistemic aims.¹²

Immodesty

Your own epistemic standard and the belief states that it recommends will always seem to you to do a better job of balancing the twin epistemic aims than any other epistemic standard and its corresponding belief states.

We can think of the immodesty assumption as a thesis about what makes some epistemic standard *yours*. As we've seen, an epistemic standard is a way of balancing the twin epistemic aims of believing truly and not believing falsely. If we want to pair this definition with standards permissivism to generate verdicts about which belief states are rational for you, we need to know which epistemic standard is yours. Immodesty provides an answer to that question: your epistemic standard is the one that seems to you to do the best possible job of balancing the twin epistemic aims of achieving true belief and avoiding false belief.

To see the immodesty assumption in action, suppose that Jane is a standards permissivist who thinks that Elizabeth has a reasonable epistemic standard and that Elizabeth's belief about Mr. Darcy is in line with that standard. Assuming that it is possible for her to adopt Elizabeth's standard and the belief states that it recommends, she couldn't justify an aversion to doing so on the grounds that she would have impermissible beliefs after the transformation—since, by hypothesis, Jane actually doesn't think that this is true. But nonetheless, because of immodesty, Jane takes herself to have a reason to resist replacing her own standards with Elizabeth's. She might describe this reason by saying: "What I care about, in the first instance, is having beliefs that strike the right balance between the aim of believing the true and the aim of not believing the false; and if I adopted Elizabeth's standards, I wouldn't have such beliefs. It is neither here nor there that I would also think, mistakenly, that my new ways of believing did a better job of balancing the twin epistemic aims than my old ways of believing." This, according to the standards permissivist, articulates Jane's distinctively epistemic objection to arbitrary switching.

¹² This response to the arbitrariness objection is developed in Schoenfield (2012) pp. 201–2 and Schoenfield (2022) pp. 281–4.

3.3 Immodesty and valuing

I think it is worth stressing something here that, to my knowledge, has not been explicitly said in existing debates about standards and arbitrariness: the immodesty assumption implies that you do not merely *prefer* or *incline* towards the way of balancing the twin epistemic aims that corresponds to your epistemic standard. This is because preferences, inclinations, tastes, and the like do not have the structural feature that the immodesty assumption attributes to our attitudes towards our own epistemic standards. Mere desires, unlike the attitudes that we are supposed to take towards our own epistemic standards, are not invariably attended by the judgment that we have reason to maintain those desires or promote their objects.¹³

Here is an example to illustrate the point: Suppose that Elizabeth inclines towards intellectual adventurousness. In other words, she is a bit more drawn to true belief and a bit less averse to false belief than most people, and her beliefs reflect these preferences. Unfortunately, Elizabeth doesn't actually endorse these preferences. She thinks of her own intellectual adventurousness as *rashness*—she thinks that it doesn't reflect the real values of the twin epistemic aims. Just as she is reflecting ruefully on this, along comes a peddler with a magic pill that will make her stop believing in Mr. Darcy's guilt and simultaneously change her epistemic inclinations so that they align with her considered value judgments. Given Elizabeth's attitudes towards her current epistemic inclinations, would it make sense for her to refuse the peddler's pill on the grounds that taking it would induce states that are not recommended by these inclinations? Clearly not.

As this example helps us to see, it is perfectly coherent to simultaneously prefer some way of balancing the twin epistemic ends and not take yourself to have reason to maintain that preference or pursue that balance. But insofar as you have some epistemic standard, you must take yourself to have reason to maintain this standard and occupy the belief states that they recommend; this is just what the immodesty assumption says. So, your attitude towards the balance corresponding to your own epistemic standard is not a mere desire. The structural property that the immodesty assumption attributes to our attitudes towards epistemic standards is a distinguishing mark of the class of attitudes that are sometimes called *valuing attitudes*.¹⁴

4 The problem with the standards-based response

We've just seen why it is crucial for the standards-based solution that your epistemic standards correspond to ways of *valuing*, rather than merely desiring, the twin epistemic goals of believing truly and not believing falsely. What I'd like to do now is raise a question about the valuing attitudes that we are supposed to take towards

¹³ Scanlon argues for this point persuasively in Scanlon (2000) p. 43.

¹⁴ Here is how Watson puts the point in his famous discussion of the difference between merely desiring and valuing: "Part of what it means to value some activities in this way is this: we judge that to cease to have such appetites is to lose something of worth. The judgment here is not merely that, if someone has these appetites, it is worth while (*ceteris paribus*) for him to indulge them. The judgment is rather that it is of value to have and (having them) to indulge these appetites" ((Watson, 1975) p. 213).

our own epistemic standards and identify some serious difficulties that the standards permissivist will face in trying to answer it.

The question can be put like this: Is the valuing attitude that a subject takes towards her own epistemic standard *categorical* or *conditional*? When you categorically value some end, you think that everyone ought to value it as you do; you judge anyone who doesn't share your attitude towards this end to be making an important mistake. When you conditionally value some end, you think that you and everyone in relevantly similar circumstances—e.g., anyone with a relevantly similar personal history, set of ground projects, or character—ought to value it as you do; you judge that *these* people would make a real mistake if they did not value the end in question. But you also think that there are plenty of people who are not relevantly similar to you, and thus have no reason to value this end as you do.

Should the standards permissivist say that our epistemic standards correspond to categorical or conditional valuing attitudes? It turns out that either way of going presents deep challenges.

4.1 The categorical option

Suppose first that to have an epistemic standard is to take a categorical valuing attitude towards that standard. On this hypothesis, if a belief state B is epistemically permitted for someone in light of evidence E —and is therefore uniquely recommended by her epistemic standard in light of E —she will judge that everyone who fails to occupy B in light of E is making a distinctively epistemic mistake, one that she can roughly describe by saying that they are not doing the best possible job of balancing the epistemic aims of approaching truth and avoiding falsity.

A worry about this hypothesis is that it attributes to everyone (everyone, at least, who has epistemic standards and is therefore a candidate for holding rational beliefs) a view that sounds very much like uniqueness. For we can easily imagine a uniquer who says: “My view is that, for any given body of evidence, there is a single doxastic response that reflects appropriate concern for getting at truths and avoiding falsehoods. Anyone who likes may have the label *epistemically permitted*, and apply it to someone's belief state only if that belief state is recommended by the epistemic standard that she happens to have. What I am fundamentally committed to is the idea that there is a single, epistemically *best* way of responding to any given body of evidence. So really, I agree with anyone that the categorical standards permissivist would count as a rational believer.” But if this is right, then the standards permissivist who goes in for the categorical view is essentially embracing a kind of widespread error theory; she attributes to everyone (or everyone who is a candidate for rational beliefs, at least) faith in a kind of uniqueness. That option was always available to the permissivist as an explanation of why we are averse to changing our beliefs through arational interventions, and if she had wanted to take it, there would have been no need to go through the rigamarole of the standards-based response to the arbitrariness challenge. That response was supposed to rationalize our aversion to arbitrary switching *without* positing a wide-spread faith in uniqueness.

Note that this worry is not quite the same as the familiar worry that relativist views, when fully spelled out, seem incoherent or threaten to collapse into a kind of non-relativism.¹⁵ There is nothing incoherent about pairing standards permissivism with the claim that people categorically value their own epistemic standards, nor does standards permissivism somehow become a kind of uniqueness when paired with this claim. The reason the categorical view is troubling for the standards permissivist is that it seems to have an implication that renders her solution to the arbitrariness challenge no better than a simpler alternative, one that explains our aversion to arbitrary switching by appealing to a widespread, though mistaken, faith in uniqueness. So, to stave off this worry, the standards permissivist who favors the categorical view needn't explain why standards permissivism is compatible with that view; rather, she needs to explain why someone with epistemic standards is not, according to the categorical view, thereby committed to uniqueness.

One way of doing this involves recasting the uniqueness/permissivism debate in semantic terms, as some contextualists and relativists do for the debates in which they are engaged. For instance, she could say something like this: "Properly spelled out, my view is that a claim of the form *B is epistemically rational* is only true or false relative to an epistemic standard. My opponent, the uniquer, denies this. This is the real difference between our views. So, the fact that everyone with epistemic standards will say things like *There is only one rational response for each body of evidence* and *There is a single best epistemic standard for everyone* does not show that they disagree with my view—because my view is not a first-order theory of epistemic rationality or epistemic value at all. It is a thesis about which parameters need to be given a value before attributions of epistemic rationality can be true or false."¹⁶

The version of standards permissivism just described might have a lot going for it, but I take it to be a substantial departure from standards permissivism as it was originally conceived. The original standards permissivism was not a semantic thesis about the meaning of the term *epistemically rational*; it was a metaphysical thesis about the grounds of epistemic rationality. Its defenders explicitly say things that the semantic permissivist just described has rejected in the hopes of avoiding disagreement with everyone who categorically values their own epistemic standards. For instance:

[I]t is possible for different rational individuals to have the same evidence, but different attitudes in response ((Pettigrew, 2022) p. 17).

The uniquely reasonable response for you is to suspend judgment about whether *p*, and the uniquely reasonable response for me is to believe *p* ((Kelly, 2014) p. 304).

[W]hat one ought to believe depends, in part, on what epistemic standards one has. On this view, if two people with the same evidence reasonably have different

¹⁵ Something like this worry is spelled out for standards permissivism by Simpson ((Simpson, 2017) pp. 529–31). Horowitz highlights a concern that is closer to the one I articulate here, although hers is not framed as a criticism of the standards-based response to the arbitrariness challenge. Rather, it highlights a worry about how moderate permissivists (a group to which standards permissivists belong) are going to explain the value of epistemic rationality ((Horowitz, 2014) pp. 49–53).

¹⁶ Some of the most important and well-developed versions of this general relativist strategy are found in Gibbard (2003) and MacFarlane (2014).

opinions about whether p , it is because these people have each adopted a different set of reasonable epistemic standards ((Schoenfield, 2012) p. 199).

Assuming that the standards permissivist who embraces the categorical option wants to continue to say such things, she cannot rely on the familiar contextualist/relativist tool of “semantic ascent” to escape the worry that valuing one’s own epistemic standards categorically is incompatible with accepting her view. Rather, she must (1) find a way of cashing out the notion of epistemic rationality that allows someone who categorically values his own epistemic standard to say that he and another person with whom he disagrees are both perfectly rational in virtue of how each balances the twin epistemic aims, even though only *his* belief states reflect the *best* way of balancing the twin epistemic aims; and (2) she must convince us that this kind of epistemic rationality is not only an important normative standard, but the very one about which she and the uniquer disagree. I don’t claim that there is no satisfying way to do these things. I just point out that it seems like non-trivial project, and may involve coming unmoored from the conceptual anchors that fix the notion of *epistemic permissibility* for most parties to the debate.

4.2 The conditional option

If the standards permissivist wants to avoid the difficulties that she faces if she supposes that we categorically value our own epistemic standards, she could try saying that we conditionally value our own epistemic standards. On this hypothesis, you judge that *you* (and people relevantly like you) would be making a real mistake if you were to occupy belief states other than the ones recommended by the epistemic standard that you presently have. But you also think that other people in your evidential position could have belief states that are not recommended by this epistemic standard without making any kind of significant mistake.

The problem with this proposal, in a nutshell, is that conditional valuing attitudes call out for justification in terms of facts about the valuer’s circumstances; and it is not clear how such a justification would go in the case of the conditional valuing attitudes that, as we are now supposing, you take towards your own epistemic standards. To expand: When you conditionally value something, you judge that only those in circumstances relevantly similar to yours have reason to value that thing as you do. For such an attitude to be reasonable, there must be some explanation of why only those people have reason to value the thing in question, and this explanation must be in principle accessible to you. In other words, there must be some intelligible connection between the distinguishing features of the group who purportedly have reason to value the relevant thing, and the nature of the reason to value that thing. If I think, for instance, that only people with brown hair have reason to care about their friends, this is unreasonable. Why? Because there is no intelligible connection between having brown hair and having reasons to care about one’s friends that could explain why only people with the former have the latter.

It seems to me that in order for your conditional valuing attitude towards an epistemic standard to vindicate your aversion to arbitrary switching, that valuing attitude must itself be reasonable. And if it is to be reasonable, I’ve been suggesting, there must

be a good explanation of why only you and people relevantly like you have reason to value the twin epistemic goals in precisely the way that corresponds to that standard. If all this is right, then a proponent of the standards-based solution who endorses the conditional view is on the hook for providing such an explanation.

But what could this explanation be? For inspiration, the standards permissivist might look to the kinds of explanations that we usually give when we justify conditional valuing attitudes. If you think that it would be best for you to spend an afternoon with so-and-so, but you acknowledge that many other people have no reason to do the same, you can justify your position by appealing to your friendship with the person in question. If you think that being a philosopher is better for you than becoming a stay-at-home parent, but you also think that many other people have reason to make the opposite choice, you can justify your position by saying that being a philosopher is part of what *your* life is about. Similarly, if you would rather maintain your current personality than manipulate yourself into a new one, but you deny that your own personality is objectively the best in any sense, you can explain yourself by saying that you would no longer be *you* if your personality changed drastically enough.

Of particular interest for our purposes are these last two examples, because we are sometimes strongly identified with our intellectual tendencies and commitments; in some cases we view them as partly constitutive of what our lives are about, or part of what makes us *us*. In light of this, it is perfectly intelligible why you might refuse to induce *many* belief states that are *substantially* different from your current ones, and why you might be particularly opposed to giving up certain moral, political, and religious beliefs. Such changes would, in an elusive but nonetheless recognizable sense, turn you into a different person.¹⁷ But, recall: In trying to respond to the arbitrariness challenge, the standards permissivist is not just looking for a reason to avoid substantial, global changes in one's doxastic profile. She is looking for a distinctively epistemic vindication of the aversion to arbitrary switching that applies in every case, even if it is a case in which the relevant belief states concern some unimportant subject matter and would change in only a minor, local way. And we cannot illuminate the nature of this reason by appealing to intellectual character, because it is simply not plausible that any small change in your doxastic profile would constitute a betrayal of who you are or make you a different person in the relevant sense.

The standards permissivist could bite the bullet here and insist that, despite the *prima facie* implausibility of this idea, any change in your epistemic standard, no matter how minor, would turn you into a different person. Alternatively, she could try saying that your valuing attitude towards your own epistemic standard doesn't need to be reasonable in order to provide the basis of a vindicating explanation of why you resist arbitrary switching. But if she pursues either of these options, again she erases the advantage that her standards-based solution to the arbitrariness challenge was supposed to have over simpler alternatives. It was always open to the non-standards

¹⁷ As I read her, Callahan relies on something like this idea in crafting her response to the arbitrariness challenge: “[A]gents have the frameworks they have because of the choices they’ve made about who to be—about the values, methods, and expectations they shall employ as inquirers and opinion-formers. An agent’s framework is thus hers in a deeper sense than merely being the framework she happens to hold or the framework associated with her. One’s framework at any point in time is partly a function of one’s choices; it is an intellectual aspect of the self one is shaping” ((Callahan, 2021), p. 552).

permissivist to explain her aversion to arbitrary switching by saying: “Being agnostic in exactly these and these evidential circumstances is part of my intellectual personality!” Granted, this does not seem very plausible; but identifying the very odd commitments expressed by such a claim with epistemic standards does not seem to make it any more plausible. Likewise, it was always open to the non-standards permissivist to defend her aversion to arbitrary switching by saying: “Look, I know it’s not reasonable, but I just think it would be a mistake to change my belief!” This doesn’t seem like the kind of full-throated vindication that the arbitrariness challenge was asking for; but again, it is unclear how adding epistemic standards to the picture is going to make matters better.

So, insofar as the standards permissivist wants to defend the superiority of her own response to the arbitrariness challenge and simultaneously embrace the conditional view, what she needs to do is (1) offer a vindication of your aversion to arbitrary switching that includes an explanation of why only you and others relevantly like you have reason to value your epistemic standard, and (2) make sure that an equally good analog of this explanation is not available to the non-standards permissivist who wants to vindicate an aversion to arbitrary switching. Again, I am not claiming that this task is impossible to complete. But the prospects for completing it look rather dim.

5 An alternative response to the arbitrariness challenge

Let’s take stock. The standards-based response to the arbitrariness challenge tries to vindicate our aversion to arbitrary switching by positing that (1) which belief states are rational for you depends on your way of balancing the twin epistemic goals, and (2) your way of balancing the twin epistemic goals will always strike you as better than any alternative. But what we’ve seen is that this proposal raises as many question as it answers about the rational basis of our aversion to arbitrary switching—and it is unclear how we are going to answer these questions without erasing the appeal of the standards-based solution. So, the standards-based solution is incomplete at best, and unworkable at worst. Given all this, I think, we should ask whether there might be some other way to deal with the arbitrariness challenge. In this section, I will sketch one possibility.¹⁸

5.1 The one-thought-too-many insight

The arbitrariness challenge asks the permissivist to identify a distinctively epistemic vindication of our aversion to arbitrary switching. One natural idea is that our objection to arbitrary switching is rooted in something like our *concern for truth*. Indeed, this is where some standards permissivists start out in devising their response to the

¹⁸ For alternative ways of addressing versions of the arbitrariness objection that don’t appeal to standards, see Meacham (2013) and Meacham (2021).

arbitrariness challenge.¹⁹ But in its particular development of this idea, the standards-based response reveals a certain way of understanding what it is to have a concern for truth—and there might be alternative understandings that deserve our attention.

On the standards permissivist's explication of concern for truth, to have a concern for truth is essentially to value the ends of believing truly and not believing falsely; and, different ways of having a concern for truth are different ways of assessing the relative importance of these two ends. This is a consequentialist conception of concern for truth, according to which having a concern for truth is a matter of valuing certain outcomes.

But not all of our concerns have this kind of consequentialist structure. Consider, for instance, love of a person. If we want to understand what it is to love a person, we hobble ourselves if we assume at the outset that any good answer to this question will take the form, "When you love a person, you value outcomes of type *F*." To illustrate the point, here is one important feature of love that could not be captured by such an answer: when you love someone, certain kinds of considerations are silenced in your deliberations concerning her.

Bernard Williams illustrates this with his famous example of the husband who thinks "one thought too many" as he sets out to rescue his wife.²⁰ Faced with a choice between saving his wife and a stranger, the husband thinks to himself, "In this circumstance, the moral considerations permit and perhaps even demand that I give preferential treatment to my wife. So, here I go!" As Williams points out, there is something quite chilling about the psychological portrait just sketched; the man described is one that many of us would not want for a partner. He seems not to love his wife in the way that we want to be loved by those closest to us.

Williams' brief discussion of this example touched off a massive literature, and this is not the place to critically examine all of the lessons about love that philosophers have extracted from it. I'll just highlight what I take to be the most important of these lessons for my purposes: some of our deepest commitments involve a deliberative insensitivity to certain kinds of considerations. For instance, to the extent that the husband in Williams' example loves his wife,

...the sight or sound of his wife in danger would—in these circumstances—so fill up his consciousness that it leaves no room for thought or care about morality or the impartial point of view. But that is not all...[the] husband [would also] be comfortable with these deliberative and motivational dispositions...Recognizing that he would respond to the rescuing situation in this way, he does not feel the need, or perhaps even any interest, in finding out counterfactually whether his response would be permitted by one who adopts an impartial point of view ((Wolf, 2012) p. 84).

The loving husband, in other words, is someone who feels immediately and irresistibly compelled by his wife's distress to save her, and is untroubled by—perhaps positively grateful for—the implication that in certain ways and in certain circumstances, he is unresponsive to the full range of considerations that bear on what would

¹⁹ I am thinking in particular of Schoenfield (2022) p. 281.

²⁰ This discussion is found in the last few pages of Williams (1981).

be moral, prudent, or wise. The more general point here is that loving someone involves occupying a deliberative perspective in which certain considerations are sometimes blocked from view or rendered wholly unmoving, and from which one endorses and perhaps even celebrates this insensitivity.

Notice that this characterization of love is not about the outcomes that the loving person desires or values. I think that the loving husband does necessarily value certain outcomes, like the ones in which his wife flourishes as an individual and ones in which their relationship is healthy, joyful, and enriching. But one central lesson of Williams' example is that valuing such outcomes is not all there is to love, for the moral husband, too, values these outcomes. What sets the loving husband apart is not a set of ends, but a deliberative point of view that involves insensitivity to certain kinds of considerations.

5.2 Concern for truth as involving a deliberative insensitivity

I'd like to suggest that concern for truth is like love of a person in that it involves occupying a somewhat blinkered perspective. Just as the loving husband is sometimes unmoved by the full range of considerations that bear on the morality and prudence of his actions, someone who cares about the truth is sometimes unmoved by the full range of considerations that bear on the morality and prudence of her beliefs.

Think back to the person contemplating Pascal's wager whom we encountered when reviewing the arbitrariness challenge. She feels compelled towards agnosticism by the evidence—i.e., the considerations bearing on the truth of theism—but then it is pointed out to her that agnosticism is not the most prudent of her options. Strikingly, she can't really bring herself to care about this. She doesn't feel any inclination to go to church or pursue other projects of self-indoctrination in order to get herself to believe that God exists. And far from viewing this insensitivity to the prudential considerations as a regrettable, akratic deficiency, she is oddly proud of it. She is happy to be the sort of person whose consciousness is "filled up" and pushed irresistibly in a certain direction by the evidential considerations as she deliberates about what to believe, so that no room is left for the considerations that bear on the prudence of the various doxastic options. Why does she feel this way?

One hypothesis—the one favored by our standards permissivist—is that the stalwart agnostic's attitudes are best explained by some ends that she has (specifically, the ends of having true beliefs and not having false beliefs, given some specified way of balancing those two ends). What I'd like to suggest instead is that her aversion to arbitrary switching is explained by a different aspect of her concern for truth: a blinkered deliberative perspective in which she is unmoved by the considerations bearing on the morality and prudence of her belief states, and from which she approves of this insensitivity. What sorts of considerations *is* she sensitive to in this blinkered deliberative standpoint that is characteristic of a concern for truth? The considerations bearing on the truth of the propositions that she contemplates—in other words, the evidence.

Why think that having a concern for truth involves occupying the deliberative standpoint just described? This is a big question, and one that I don't have the space to thoroughly investigate here. But I can say a few things that go some ways towards

answering the question. First, my proposal about concern for truth creates an appealing symmetry between concern for truth and other deep commitments like love, given how these latter commitments are understood by independently motivated views. Second, it promises to help us explain much more than our attitudes towards arbitrary switching. I have a hunch that our aversion to wishful thinking, for instance, can also be partly explained by appealing to the deliberative standpoint characteristic of a concern for truth. My reasons for this hunch are basically these: Wishful thinking does not necessarily lead to beliefs that are out of line with the evidence,²¹ nor does it necessarily lead the believer to doxastic states that, from her own perspective, do a sub-optimal job of balancing the ends of believing truly and not believing falsely. What it *does* necessarily involve is occupying a deliberative perspective in which one's beliefs are influenced by non-evidential considerations. If I am right that this deliberative standpoint is incompatible with a concern for truth, we will have in hand a neat, distinctively epistemic explanation of our aversion to wishful thinking, and one that unifies it with nearby phenomena like our aversion to arbitrary switching.

5.3 A solution to arbitrariness for all—no standards required

There is much more to do to defend the idea that having a concern for truth involves occupying a blinkered deliberative standpoint. This is work that I (and hopefully others as well) may take up in the future. But for now, what I'd like to stress is that *if* this idea is right, it provides us with a very ecumenical response to the arbitrariness challenge, one that is available to the standards permissivist and non-standards permissivist alike.

Recall: The arbitrariness challenge demands that we vindicate our aversion to arbitrary switching in a distinctively epistemic way, and without appealing to an interest in having epistemically rational beliefs. Our standards permissivist tried to do this by appealing to other ends, namely, the ends of having true beliefs and not having false beliefs, given some specified way of balancing those two ends. What I'm suggesting now, by contrast, is that we give up on ends-based explanations altogether and appeal instead to facts about the blinkered deliberative standpoint characteristic of having a concern for truth. On my view, these facts can rationalize our distinctively epistemic objection to arbitrary switching, just as analogous facts about the deliberative standpoint of love rationalize the motivations of the husband who rushes in to save his wife without giving morality a single thought.

The rationalizing story goes something like this: Suppose that you occupy the blinkered deliberative perspective that, on the present proposal, is partly constitutive of having a concern for truth. This means that you are insensitive to considerations bearing on the morality or prudence of your own belief states. So, if you have arrived at some belief state, the non-evidential considerations that might motivate someone else to arbitrarily switch out of it—e.g., the prudential benefits of switching to an alternative state—will leave you cold. Indeed, the idea of responding to such considerations will strike you as very distasteful, perhaps even alien.²²

²¹ See Ellis (2022) for a nice discussion of this point.

²² Notice that in this story, you do not reject arbitrary switching because you see something wrong with arbitrary switching *per se*. Rather, you reject arbitrary switching because the kinds of considerations that

To embrace this kind of response to the arbitrariness challenge, you do not need to be a standards permissivist; you can be a plain old permissivist.²³ In fact, you do not need to endorse *any* thesis about epistemic rationality at all. What you do need to do is accept that the domain of distinctively epistemic normativity—roughly, the normativity that binds us in our capacity as people who care about the truth—cannot be fully explained in terms of its relationship to outcomes. In other words, you must accept that some of our distinctively epistemic commitments are not merely commitments to ends.

6 Implications for debates about transformative choice

So far, we have been immersed in a debate about permissivism, one centered on the question of whether permissivism has the resources to vindicate our aversion to arbitrary switching in a distinctively epistemic way. The standards permissivist's answer to this question appealed to the hypothesis that what is epistemically rational for you to believe depends in part on your way of valuing the twin epistemic goals, and that arbitrary switching, insofar as it would lead you to have epistemically rational beliefs, would also lead you to change your way of valuing the epistemic goals. In the terminology that Paul introduced and that I referenced in the introduction, the standards permissivist treats arbitrary switching that yields rational belief as involving a kind of *personally transformative choice*—one that will change what the subject cares about at the deepest level—and this treatment is a key element of her preferred response to the arbitrariness challenge. I raised some worries for that response and proposed an alternative that is available to the standards and non-standards permissivist alike.

Now I'd like to surface from this debate about permissivism and arbitrariness in order to highlight an implication of the foregoing discussions for broader questions

Footnote 22 continued

the arbitrary switcher responds to strike you as wholly unmoving.

Because my story has this feature, I can see someone worrying that it does not provide a fully general explanation of why we reject arbitrary switching in every case. Here is one way to articulate the worry: "The present proposal explains why we are reluctant to switch out of our belief states for moral or prudential reasons. But it doesn't explain why we would reject, say, popping a pill to switch from one permissible belief state to another *for no reason at all*—i.e., when there is nothing to be gained from making the switch." I want to simply agree with the main contention here. My proposal *doesn't* explain why you would stick with your own belief state when there is absolutely nothing to be gained from making a switch. But I take it that this fact is not one that requires special explanation, for it is simply a specific instance of a much more general phenomenon: whenever there is no special reason to change course, you stay the course. For instance, if you have selected Colgate and there is no reason to switch to Crest, you stick with Colgate. If you have selected bushel 1 and there is no reason to switch to bushel 2, you stick with bushel 1. And if you are agnostic and there is no special reason to switch to belief, you stick with agnosticism.

The arbitrariness challenge gets its teeth from the fact that in some cases, there *are* clear benefits to switching out of the permissible belief state that you currently occupy. In these cases, we need a special explanation of why you would refuse to switch—and that is what my proposal aims to provide.

²³ This is not to say that the response I am endorsing here is logically *incompatible* with standards permissivism. But if it is successful, it does remove one important motivation for taking on the explanatory burdens and costs in simplicity associated with being a standards permissivist—because it offers a response to the arbitrariness challenge that is available to a permissivist who endorses the much simpler, "naive" form of the view. So, the proposal that I have sketched in this section does somewhat undermine standards permissivism, even though it is not incompatible with it.

about personally transformative choice. In a nutshell, that implication can be articulated as follows: If Williams and his sympathizers are right in thinking that not all of our deepest commitments are (just) commitments to ends, then one of the original ambitions of the literature on transformative experience is fundamentally misguided. This ambition is to come up with unified answers to very general questions about how to make rational transformative choices, and articulate these answers within the framework of decision theory. For instance, here is Paul voicing some of the questions that interest her:

In contexts of transformative choice, how are we to make decisions within the constraints of deliberative decision theory? How are we to determine and follow the relevant diachronic rational norms? ((Paul, 2015) p. 801)

[H]ow are we to weight our local utility functions over time when framing and contemplating the possibility of transformative change...? ((Paul, 2015) p. 806)

But if not all of our commitments can be wholly understood as commitments to ends, then we shouldn't expect unified answers to very general questions like, "How should you decide whether to change one of your deepest commitments?"; we shouldn't expect a single set of diachronic rational norms to govern all transformations of commitment; and we certainly shouldn't expect all norms governing transformations of commitment to be translatable into decision-theoretic terms. Where some commitments are concerned—commitments that take the form of valuing certain ends seem to be the most promising candidates—deliberative decision theory might prove to be a good framework in which to articulate rational norms of transformation.²⁴ But where other kinds of commitments are concerned, it seems like deliberative decision theory really isn't going to help us think about whether to maintain or abandon those commitments.

Commitments like love and (on my view) concern for truth are cases in point. These commitments involve occupying a blinkered deliberative perspective in which certain considerations are silenced, and from which you endorse this silencing. Plausibly, the silenced considerations often include considerations bearing on the prudence of maintaining the commitment in question. (If it would be disconcerting for the husband to save his wife in part because he has decided morality permits it, how much more disconcerting would be if he saved his wife in part because he judged continuing to love her would be prudent all-things-considered, or prudent with respect to some narrower goal!) But to reason about a choice in a decision-theoretic way is, fundamentally, to reason about it from the perspective of prudence. It is to ask how well each of the options available to you would advance your ends, or some narrower set of those ends. So, insofar as you are using decision-theoretic norms to answer the question of whether to maintain or abandon your love for someone, your love for them is already significantly corroded. You have already lost to a large extent the commitment that you are considering giving up, because having that commitment is partially constituted by

²⁴ Although I do have some reservations about even this idea, because decision theory tells us how to act in light of some ends; it doesn't tell us how to choose our ends. For a nice expression of this point, see Srinivasan (2015).

occupying a deliberative perspective from which the decision-theoretic reasoning you are currently pursuing would leave you completely cold.

My suspicion is that many of our deepest commitments are structured like love of a person and concern for truth. If this is right, then in many cases there is very little to be said in decision-theoretic terms that can help us decide when to give up a deep commitment. Again, this is because insofar as you have one of these love-like commitments, certain questions—which notably include *Would changing this commitment be prudent?*—strike you as completely irrelevant, and the considerations that bear on them, wholly unpersuasive; so, as soon as you take this question seriously, it becomes moot. The upshot is that, if we want to provide useful philosophical advice about making certain transformative choices, we will need to move beyond questions about decision theory and the species of rationality that it aims to capture.

Declarations

Conflict of interest The authors have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Austen, J. (1918). *Pride and prejudice*. Scribner.
- Callahan, L. F. (2021). Epistemic existentialism. *Episteme*, 18, 539–554.
- Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, 116(2), 187–217.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46(3), 279–295.
- Ellis, J. (2022). Motivated reasoning and the ethics of belief. *Philosophy Compass*, 17, 6.
- Friedman, J. (2020). The epistemic and the zetetic. *The Philosophical Review*, 129(4), 501–536.
- Gibbard, A. (2003). *Thinking how to live*. Harvard University Press.
- Greco, D., & Hedden, B. (2016). Uniqueness and metaepistemology. *Journal of Philosophy*, 113(8), 365–395.
- Hieronymi, P. (2004). Controlling attitudes. *Pacific Philosophical Quarterly*, 87, 45–74.
- Hieronymi, P. (2005). The wrong kind of reason. *Journal of Philosophy*, 102(9), 437–457.
- Horowitz, S. (2014). Immoderately rational. *Philosophical Studies*, 167, 41–56.
- Horowitz, S. (2017). Epistemic value and the Jamesian goals. In J. Dunn & K. Ahlstrom-Vij (Eds.), *Epistemic Consequentialism*. Oxford University Press.
- Kelly, T. (2014). Evidence can be permissive. In M. Setup, J. Turri, & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 298–310). Wiley.
- MacFarlane, John. (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford University Press.
- Meacham, C. (2013). Impermissive Bayesianism. *Erkenntnis*, 76(6), 1185–1217.
- Meacham, C. (2021). Arbitrariness and uniqueness. *Pacific Philosophical Quarterly*, 102(4), 665–685.
- Paul, L. A. (2014). *Transformative experience*. Oxford University Press.
- Paul, L. A. (2015). Transformative experience: Replies to pettigrew, barnes, and campbell. *Philosophy and Phenomenological Research*, 91(3), 794–813.

- Pettigrew, R. (2022). *Epistemic risk and the demands of rationality*. Oxford University Press.
- Rosen, G. (2001). Nominalism, naturalism, epistemic relativism. *Philosophical Perspectives*, 15, 69–91.
- Scanlon, T. M. (2000). *What we owe to each other*. Harvard University Press.
- Schoenfield, M. (2012). Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Nous*, 48(2), 193–218.
- Schoenfield, M. (2022). Meditations on beliefs formed arbitrarily. In *Oxford Studies in Epistemology*, Vol. 7, Chap. 10, pp. 278–305. Oxford University Press.
- Simpson, R. M. (2017). Permissivism and the arbitrariness objection. *Episteme*, 14(4), 519–538.
- Singer, D. (forthcoming). What epistemic reasons are for: Against the belief-sandwich distinction. In Dunaway, B., Plunkett, D, (eds.) *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*. Maize Books.
- Srinivasan, A. (2015). All the same. *The Times Literary Supplement*, pp. 1–3.
- Watson, G. (1975). Free agency. *Journal of Philosophy*, 72, 205–220.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19, 445–459.
- Williams, B. (1970) Deciding to believe. In *Problems of the Self*, pp. 136–151. Cambridge University Press.
- Williams, B. (1981). Persons, character, and morality. In J. Rachels (Ed.), *Moral Luck*. Cambridge University Press.
- Wolf, S. (2012). One thought too many: Love, morality, and the ordering of commitment. In U. Heuer & G. Lang (Eds.), *Luck, Value, and Commitment: Themes from the Ethics of Bernard Williams* (pp. 71–92). Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.