

Epistemic Standards and Minimally Responsive Beliefs

Federico Bongiorno

May 14, 2025

Penultimate version. Forthcoming in *Philosophy and Phenomenological Research*

Abstract

It is widely held that beliefs respond to evidence. However, it is not easy to make precise exactly in which sense beliefs are so responsive. In this paper, I develop and defend a novel minimalist account of evidence-responsiveness. I argue that in order for an attitude A to count as minimally responsive to the evidence, e , the attitude holder must follow at least one epistemic standard which appropriately explains her holding A given e . The account given here allows that attitudes may still show responsiveness in this minimal sense, even in cases where we ignore evidence against them or when we lack the capacity to revise them. It also offers a systematic way of demarcating the boundaries between beliefs and other attitudes like imaginings and acceptances. While beliefs are not the sole attitudes whose evidence relations are bound by epistemic standards, beliefs and non-beliefs undergo different forms of regulation.

1 Preliminaries

There is widespread agreement that part of what it is for something to be a belief is for it to be responsive to evidence. Typically, this amounts to either (or both) of two claims: that beliefs are disposed to be formed on the basis of, and extinguished if they conflict with, evidence in possession of the subject.¹ An attitude that is not so disposed is not a belief, or at best it is unrecognisable as such. A primary motivation for this view comes from the thought that beliefs are the output of a mechanism that is designed to get at truth (Shah & Velleman, 2005; Velleman, 2000; Millikan, 1984). Beliefs are responsive to evidence in that they are regulated for truth, and evidence is a means for attaining truth (Shah & Velleman, 2005, p. 500).

Empirically speaking, it is doubtful that we are endowed with an across-the-board disposition to adjust beliefs in light of evidence. There are at least two problematic sets of empirical results. The first concerns a phenomenon sometimes termed ‘the illusory truth effect’ (Hasher et

¹ Those who endorse both claims include: Shah and Velleman, 2005; Egan, 2008; Adler, 2002. Those who endorse the latter include: Currie and Ravenscroft, 2002; Levy, 2015; Van Leeuwen, 2014; Gendler 2008; Helton, 2020.

al., 1977). Research has demonstrated a robust correlation between message repetition and a proneness to believe the repeated message even when one knows it is false or even implausible (Fazio et al., 2019; Lacassagne et al., 2022). The greater the exposure, the more likely one considers it to be, regardless of the evidence available. The second is a phenomenon known as ‘belief perseverance’, whereby people are systematically inclined to cling to their beliefs despite counterevidence.²

If beliefs are designed, or at any rate, disposed, to be regulated by evidence, why do they fail to be so regulated as often as they do? One could insist—quite reasonably, perhaps—that design may well correlate with occasional or even frequent malfunctioning. Most mechanisms are fallible and error-prone. So, it may happen that some token beliefs fall short of manifesting evidence-responsiveness, since while they are, by design, responsive to evidence, they need not to successfully perform that function at all times. However, this leaves room for the possibility that beliefs are regulated by mechanisms for something other than evidence, and thus not—at least not regularly or even often—evidence-responsive. For even if we grant by way of concession that evidence-responsiveness ‘doesn’t require beliefs to be governed by truth-seeking mechanisms alone’ (Velleman, 2000, p. 254), it is unclear why we should suppose that beliefs remain evidence-responsive even where they are ill-grounded or evidence-resistant (Flores, forthcoming).

Alternatively, one might simply deny that there are any such things as ill-grounded or evidence-resistant *beliefs* (Currie and Ravenscroft, 2002, p. 15; Gendler, 2008, p. 566). That is, one might disqualify attitudes from counting as genuine beliefs if they are not manifestly evidence-responsive. Yet this response falls prey to what Glüer and Wikforss (2013) call the ‘extension problem’: we are left with the counterintuitive result that we have a few beliefs, for many beliefs that we attribute in everyday life lack any discernible support and defy counterevidence (Bortolotti, 2009).³

If we want to avoid this problem without doing away with evidence-responsiveness, then our challenge is to account for the fact that beliefs can remain constitutively evidence-responsive even while being held in the absence of evidence, and despite counterevidence. One way of trying to meet the challenge is to set a minimal threshold for constitutive evidence-responsiveness that

² Among landmark studies, see Ross et al., (1975) and Anderson et al. (1980); for overviews, see Jelalian and Miller (1984) and Chan et al., (2017).

³ Some philosophers have claimed that several cases of apparent ill-groundedness and evidence-resistance are in fact cases of appropriate responsiveness to evidence. This strategy has been suggested in the case of belief polarisation, whereby people become more extreme and entrenched in their beliefs after encountering both confirming and disconfirming evidence (e.g., Westfall, 2024; Dorst, 2023; Begby, 2024), and beliefs that run contrary expert consensus (‘bad beliefs’), like beliefs that vaccines cause autism or that the threats from climate change are exaggerated (Levy, 2022). Space constraints oblige me to set these views aside. It is worth noting, though, that if they are true then the extension problem ceases to be a problem. My goal here will be to provide a different sort of argument for the same conclusion.

could meet two negative desiderata: it should make the category ‘belief’ neither too broad (incorporating attitudes other than beliefs), nor too narrow (excluding attitudes that we readily ascribe as beliefs).

There have been two recent views along these lines. The first builds on Grace Helton’s (2020) proposal of a revisability condition on belief, which says that a person counts as believing p just in case they are capable of revising it in response to counterevidence bearing against p . On this view, which has been further developed by Carolina Flores (2021a; forthcoming), even groundless and recalcitrant attitudes will fall above some threshold-level of evidence-responsiveness just by virtue of being revisable when appropriate.⁴ Another approach is taken by Marianna Bergamaschi-Ganapini (2020; forthcoming), who argues that beliefs are responsive to evidence in at least a minimal sense just so long as they react to ‘detected irrationality’ in a way intended to overcome conflict (even should the reaction in itself make the holder out to be more irrational).

My first task in this essay will be to argue that even if each of these accounts go some way towards meeting the desiderata identified earlier, they either leave issues unresolved or create new areas of concern. The problem with the former is that it divorces the criteria for being evidence-responsive from the actual processes of attitude formation and maintenance. The problem with the latter is that it says nothing about cases of belief that do not react to the evidence in the way envisaged, and overlooks that some non-doxastic attitudes do.

My second task will be to develop an alternative minimum threshold requirement for evidence-responsiveness, which I shall call *evidential minimalism*. Here’s the core idea: if some attitude A is minimally evidence-responsive, then, for some set of evidence e , the attitude owner follows an *epistemic standard* which appropriately explains her holding A given e . Exactly how all this is meant will become clearer later in the paper. For now, suffice it to say that epistemic standards, in my use of the term, are policies that fix the type and weight of evidence bearing on whether p . Note that a standard can still count as epistemic even if it does not reliably track truth, provided the person employing it takes it to be reliable. A standard which recommended believing on e because doing so makes one’s life easier would not track evidence-responsiveness in any interesting sense.

The main import of my view is to construe minimal evidence-responsiveness in a way that captures the following fact — a fact that rival views do not quite capture: beliefs may be poorly grounded, recalcitrant, or even unrevisable and still be evidence-responsive by their holders’ lights

⁴ Helton and Flores defend very similar views, and since nothing much hangs on their differences for my purposes here, I will mostly focus on Helton’s and integrate Flores’ insights when relevant.

(viz. by standards their holders take to be instrumental to getting and preserving truth). As such, my view maintains an essential connection between being a belief and aiming at truth, but cashes it out in terms of the many different epistemic standards that agents can have. I will show how this view can accommodate likely instances of beliefs that do not quite meet more demanding versions of the evidence-responsiveness threshold, like those contributed by Helton-Flores and Bergamaschi-Ganapini. I will also argue that it offers a helpful way of distinguishing beliefs from other cognitive attitudes such as acceptances and imaginings.

The roadmap is as follows. In §2 and §3, I will argue that Helton-Flores and Bergamaschi-Ganapini's proposed solutions to the extension problem do not withstand close scrutiny. In §4, I will introduce and articulate the notion of epistemic standards, which lays the foundation for evidential minimalism. In §5, I will argue that evidential minimalism can avoid the pitfalls of the other views. I will close in §6.

2 The revisability view

An ongoing theme in recent philosophy of mind and epistemology has been the extent to which beliefs are evidence-responsive. As we have already noted, if we take evidence-responsiveness in a strong sense, then we are hard-pressed to explain why we so often sustain belief despite the lack of evidence or even in the face of counterevidence. One way of trying to get around this difficulty is to deny belief status to such cases. But then another problem emerges. Recategorising attitudes that we ordinarily call beliefs as non-doxastic risks narrowing down the category belief too much. Perhaps the most obvious way out of the predicament is to relax the standard of how an attitude needs to be in order to qualify as evidence-responsive.

In this spirit, Helton (2020) has offered a view on which the evidence-responsiveness benchmark could be met independent of what beliefs actually do or are set to do; her 'revisability view of belief' (RVB) proposes that beliefs are constitutively evidence-responsive in the sense that it must be possible for subjects to revise them when they are contravened by the evidence. Helton draws here on a version of the principle 'ought implies can' according to which if it is the case that one ought rationally to revise one's beliefs in response to counterevidence (as she thinks it is), then one must be psychologically capable of doing so. Since revisability alone suffices for evidence-responsiveness, even attitudes that are not formed and revised (or designed to be formed and revised) on the basis of evidence can still count as evidence-responsive (and thus as beliefs), so long as they are revisable. Though if an attitude is such that it cannot be so revised, then it is not a belief but some other attitude in its vicinity, because beliefs are necessarily revisable.

Note that the capacity to revise one's beliefs is fallible in that it is possible for there to be an agent, S , such that: S retains the ability to Φ even when S is not Φ ing, including when S tries, but fails, to Φ . When Helton illustrates this idea, she adopts a counterfactual formulation: to say that S is able to revise an attitude A is to say that S would rationally revise A in a sufficiently broad range of circumstances in which S 's overall psychology is similar, but not identical, to its current state (2020, p. 513). An example Helton gives is a case in which we have evidence that a friend stole from us but we continue to believe that they are innocent because of the affection we have for them. Though our affection for our friend is preventing rational revision from taking hold, there are nearby worlds in which our mental states are different from our current ones, such that our affection is absent and we rationally revise our belief.

Flores (forthcoming) casts RVB more explicitly in terms of 'masking' interferences. So construed, Helton's case is one where we have the ability to rationally respond to evidence but some interfering factor (in the example, an affection) masks its successful exercise. If the impediment were removed, the ability would be successfully exercised, resulting in a rational revision of belief. Where Flores disagrees with Helton is over the conditions that must be in place for one to count as having the ability to revise. Helton seems to leave some wiggle room for the ability to fail even when all the relevant evidence is in and masks are absent. For Flores, any such failure would count against the possession of the ability.

Either way, RVB offers us an elegant paradigm to test whether a certain attitude A is a belief or something else. Quite simply, if the ability to revise A is altogether missing rather than merely unexercised, as when a mask prevents it from manifesting, then A is not constitutively evidence-responsive, and thus it is not a belief. Flores (forthcoming) applies this criterion as a ground for addressing the extension problem. Her strategy depends on two claims. The first is that familiar cases of evidence-resistant beliefs are best explained not by the subject lacking the capacity for revision, but by the presence of masking conditions in the form of defence mechanisms.

The argument goes, roughly, like this. We know from extensive psychological research that people feel uncomfortable when they encounter information that is inconsistent with their cherished beliefs — e.g., ones that affirm their own sense of identity or self-worth (e.g., Mandelbaum, 2019, p. 12). This unpleasant state of mind, which Leon Festinger called 'cognitive dissonance' (1957), motivates one to treat evidence asymmetrically: confirming information is accepted uncritically, whereas information to the contrary is overly scrutinized. Unless evidence-resistant beliefs were capable of being revised, there would be no purpose for such defence mechanisms against revision, since counterevidence would pose no threat. If so, real-world

evidence-resistance is not just compatible with, but supportive of the idea that beliefs are constitutively evidence-responsive.

Flores' second claim is about how to set apart beliefs from non-doxastic attitudes like imaginings: while both beliefs and imaginings are revisable, imaginings are likely to remain unrevised upon receiving counterevidence, even when conditions are optimal for the exercise of one's capacity to revise. Unlike beliefs, imaginings are not required by their nature to represent reality, and as such they do not constitutively require the capacity to revise in light of evidence. If they did, argues Flores, one of their primary roles would have to be given up, namely their role in representing alternatives to reality.

Now it is apparent that RVB does well by our desiderata. Evidence-responsiveness as thus conceived allows for even deeply evidence-resistant attitudes to count as beliefs, while at the same time respecting the need to demarcate beliefs from attitudes that are uncontroversially non-beliefs. Yet there is a question (so I will argue) as to whether this is at too high a price.

Suppose a given person's brain undergoes some degree of physiological change which makes her unable to revise an attitude she has. In Helton's leading example, Sheryl has a well-supported view that her neighbourhood has a weekly market on Fridays. Sometime after, she sustains a mild brain injury which selectively impairs her ability to relinquish her view that the market is on Fridays. As a result, she continues to hold onto this view even as she keeps gathering evidence that the market now operates on Sundays. Helton continues the example by describing how the injury has the retroactive effect of transforming Sheryl's belief into an *idée fixe*, which is a placeholder for a belief-like attitude that is unrevisable. The point, for Helton, is that a belief ceases to exist as soon as it ceases to be evidence-responsive, which happens as soon as its holder loses the ability to revise it.

There is nothing conceptually problematic with this claim, but it leads to unpalatable consequences. If RVB is true, then what makes an attitude evidence-responsive is not that it does behave in a certain way, but that it would behave in a certain way under the proper conditions. So attitudes can remain or cease to be evidence-responsive independent of the actual processes that form and sustain them. This will result in us having to say of attitudes that are based rationally on the available evidence that they are not evidence-responsive when actually they are.

Imagine a scenario in which Sheryl forms her view on good evidence but never encounter any reason to give it up. Indeed, as it happens, she only gains evidence to further support it. One day, something clicks inside her head and she loses her capacity to revise it. RVB would have us say that the view Sheryl holds is not evidence-responsive and thus not a belief, even though it is based on sufficient evidence and there is no evidence it is failing to respond to.

To me it is counterintuitive to suppose that that even where Sheryl responds to all of her evidence rationally, her view counts as unresponsive merely on the grounds that she would not revise it in circumstances that might never even arise. More generally, it is counterintuitive to suppose that someone cannot end up with a responsive attitude, even if they have responded to all of their evidence rationally. Helton and Flores might want to bite the bullet and accept such results. Indeed, they might argue, it is not too hard to bite the bullet here, since the bullet is part of an attractive approach to the extension problem.

I do not find this line of response very convincing. While RVB can block the implication that we have very few beliefs, it still gets the extension of belief wrong, for it misclassifies an intuitive case of evidence-responsiveness as one of unresponsiveness. Regardless, my worry is not so much whether the bullet is a hard one to bite, but rather whether it is a bullet we need or should bite. RVB lives up to its bill, but it does so at the cost of creating a disconnect between the reasons for which one forms or maintains an attitude and what makes it evidence-responsive. Whether the outcome is worth the cost will depend on whether there is an adequate response to the problem which avoids misclassification errors; that would involve tying evidence-responsiveness more closely to the way attitudes actually behaves. I will now examine one such response and the limits it encounters before providing my own.

3 Moderate traditionalism

Bergamaschi-Ganapini's view is broader and in a way more ambitious than RVB, in that its concern is not only with evidence-responsiveness, but with rationality broadly conceived. She calls it 'moderate traditionalism' (MT), since it takes a middle path between two extremes: 'strong traditionalism', according to which in order for something to be a belief, it must comply with norms of rationality, and 'revisionism', according to which belief may be possible without complying to such norms. MT holds that beliefs are at least minimally rational in the sense that they react to 'detected irrationality' by restoring consistency. Irrationality can take a variety of forms, but it is epistemic irrationality that we are interested in here. The type of epistemic irrationality that MT is concerned with is 'structural inconsistency', the internal incoherence of some system of beliefs within itself (Wedgwood, 2023, p. 61).⁵

With this in mind, we can see that MT gives us a new way of understanding evidence-responsiveness, not in terms of revisability, but rather in terms of a certain type of inconsistency-resolution: whenever evidence one acquires is discovered to run counter something else that one

⁵ In a forthcoming paper, Bergamaschi-Ganapini develops a version of MT applied specifically to epistemic irrationality, which she calls simply 'minimalism'. For simplicity's sake, I only refer to MT here.

believes, it creates cognitive dissonance (to recall, a state of discomfort that prompts an unconscious urge to restore coherence to one's belief system).⁶ An example would be a case where one believes that p , acquires some evidence e , and believes that e conclusively indicates that not- p (Bergamaschi-Ganapini, forthcoming). This suggests a test for whether an attitude can come within the scope of belief: if being confronted with counterevidence that one acquires and recognises as such causes no discomfort or drive towards coherence, then the attitude in question is not a belief.

Two clarifications before proceeding. First, there are different ways in which coherence could be restored, and not every way will improve one's epistemic position. Indeed, some will make it worse. Upon discovering evidence that not- p , one might give up one's belief in p , but one might also dismiss the conflicting evidence, explain it away by recourse to new beliefs, or even avoid it altogether. According to MT, all that matters for minimal evidence-responsiveness (and hence belief status) is the arousal of dissonance brought about by the conflict and the correlative drive to resolve it. This is independent of whether resolving the conflict draws one closer to the truth (e.g., selectively avoiding evidence clearly will not). So MT can accommodate the fact that beliefs are often held despite counterevidence. As long as they are accompanied by discomfort and a push for coherence, such cases count as beliefs.

Second, evidence-responsiveness so understood is not only necessary for belief, but it is also crucial for distinguishing belief from other cognitive attitudes such as accepting or imagining. Focusing exclusively on the latter, Bergamaschi-Ganapini argues that beliefs and imaginings can be distinguished in part because, and to the extent that, they react to counterevidence in different ways. As far as imaginings are concerned, there is no drive impelling the agent toward any sort of inconsistency-resolution (with a few exceptions, such as when these occur in single imaginative episodes). For example, when one compares incompatible courses of action, one can entertain alternative perspectives without being moved to seek coherence. If this were right, MT's conditions for evidence-responsiveness could be worked into a principle of demarcation between beliefs and non-beliefs.

Despite its strengths, MT has two limitations. The first limitation is that we do not seem to feel discomfort or react in the way MT envisages every time our beliefs are challenged by

⁶ Cognitive dissonance is the subject of extensive and ongoing research in social psychology. It has been studied in a wide assortment of experimental paradigms, and the proper interpretation of the results is still up for grabs. Since my goal here is only to introduce the main ideas in MT, I will stick close to Bergamaschi-Ganapini's own discussion, which draws largely on Festinger's original theory (1957). For a helpful recent review of the cognitive dissonance literature, see Harmon-Jones and Mills, 2019.

conflicting evidence, and so it is unclear how we should treat these kinds of cases.⁷ Suppose John believes the coffee machine is broken. A female colleague, Esther, drops in on him in his office to tell him that the coffee machine is working. He has no reason to distrust her, but since he trusts himself more than others, especially if women, he ignores her testimony, and continues to believe that the coffee machine is broken. Such a scenario is not only conceivable but (sadly) familiar. Now suppose John does not drink coffee and is indifferent to whether her colleagues drink it or not. Do we have reason to think that he would experience dissonance upon hearing Esther's testimony? It seems to me that we have no such reason.

Recall that on MT beliefs necessarily retain at least a minimal level of evidence-responsiveness, construed as a dissonance-driven reaction of the sort we have just been discussing. A counterexample to MT, then, would be a case where one believes p and that her evidence shows that not- p , yet one feels no discomfort nor motivation to resolve the inconsistency. Is John's one such case? It seems not, for it looks like John does not believe Esther's testimony. If so, Bergamaschi-Ganapini could accept that John feels no discomfort, but deny that it is troublesome to MT. Since John ignores Esther's testimony, his beliefs remain consistent, and so no discomfort ensues.

This seems right, but it leaves some matters wanting. There are many more cases like John's where a believer of p receives evidence against p but does not take it up. For example, think of people with delusions, who often receive testimonial evidence they discount because they underestimate the competence or sincerity of testifiers (Salice & Miyazono, 2020). MT leaves us in the dark about what to say in such cases; cases where evidence is available but one fails to take it up and believe what it shows. To the extent that no inconsistency arises in one's belief set, delusions of the above sort do not fall short of the mark of how MT understands minimal evidence-responsiveness. But does it follow that they are minimally responsive? If so, on what grounds? If not, why not? MT does not say (more on this below, §5.1).

An alternative line of response would be to argue that since dissonance is ubiquitous whenever evidence against one's belief is available (whether one takes it up or not), John's case is psychologically implausible. To argue in this way, however, would be to ignore potential contextual factors that may prevent counterevidence from leading to dissonance. For one, dissonance is likely to occur when the beliefs being threatened are ones a person cares about and desires to be true,

⁷ Bergamaschi-Ganapini could rectify this limitation by interpreting MT counterfactually. For, on a counterfactual interpretation, an attitude could count as evidence-responsive even when one fails to feel discomfort in the face of counterevidence, provided one would feel discomfort if certain circumstances were different than they actually are. It should be clear, though, that MT would then face the same limitation that we saw when considering RVB, of divorcing evidence-responsiveness from how attitudes actually relate to the evidence.

so-called *core* beliefs, such as that one is good, smart, consistent, etc. (Aronson, 1969; Bendana & Mandelbaum, 2021; Mandelbaum, 2019). When stakes are low, as in case like John's where *e* does not threaten core beliefs, this is far less likely to be the case. (Note that as long as one does not believe what *e* shows, the core belief that one is consistent is not at risk). Furthermore, sometimes people pre-emptively discredit evidence sources (Begby, 2021). So, when they confront contrary inputs from such sources, those inputs may hardly have any effect (Nguyen, 2020). Importantly, discredit need not be driven by need to avert future discomfort, but may (and often does) come from prejudice. In my example, John sees Esther's testimony as less credible because of a prejudice targeting her gender identity (see Fricker, 2007, Ch. 2).

A second limitation of MT is a problem of demarcation. Recall that, for Bergamaschi-Ganapini, one of the main attractions of MT consists in drawing a principled line between beliefs and closely related attitudes such as imaginings and acceptances: the latter would not exhibit the same discomfort-reaction pattern as the former. I think that this view is plausible when applied to imaginings but less clear that it holds for *acceptances*. To accept a proposition, as I use the term here, is to take it for granted as premise for reasoning and decision-making without believing it (Cohen, 1989; Bratman, 1992). For example, suppose I intend to stop by a bookstore on the way home from an early morning run. Although I have neither the belief nor sufficient evidence to believe that the bookstore is open in the morning, I may nonetheless take it for granted that it will be open because doing so simplifies my planning (adapted from Bratman, 1992, pp. 12-13).

Now suppose I also intend to pop over my neighbour Sylvia's house with some brownies I had baked, and (again, for the sake of day-planning) I accept that she will be home despite being uncertain whether she will. Unbeknownst to me, Sylvia works as a clerk in said bookstore every other day of the week, and when she does not, the bookstore remains closed. My acceptances are inconsistent, for if the bookstore is open, Sylvia is not home, and if Sylvia is home, the bookstore is closed. Upon finding out, I would realise I had committed myself to incompatible courses of action, which alone would be enough to elicit discomfort. It seems likely, then, that I would react just in the same way I would if I had found myself holding (some) inconsistent beliefs, by working to restore consistency. That means inconsistency-avoidance is not a prerogative of belief, as MT would have it, and so can hardly be used as an effective demarcation criterion between beliefs and other cognitive attitudes.

Let us take stock of this section. According to MT, beliefs (and only beliefs) are necessarily minimally evidence-responsive, with minimal evidence-responsiveness here understood as inconsistency-avoidance. As we have seen, however, a failure to manifest inconsistency-avoidance behaviour is not obviously sufficient to disqualify an attitude from being a belief. Nor is the simple

fact that one exhibits such behaviour sufficient for an attitude to be a belief. In the remainder of the paper, I will suggest that we understand minimal evidence-responsiveness in a different way.

4 Evidential minimalism

To restate my view in full:

EVIDENTIAL MINIMALISM (EM)

For any set of evidence e , an attitude A that S takes towards p is minimally evidence-responsive only if there is (at least) one epistemic standard x that S follows which appropriately explains why S holds A in light of e .

In this section, I spell out EM in some detail. In particular, I will discuss how ‘epistemic standards’, ‘following’, and ‘explaining’ should be best understood.

4.1. *Epistemic standards*

The concept of ‘epistemic standard’ lays at the heart of the debate over Uniqueness and Permissivism in epistemology (e.g., White, 2005; Feldman, 2010). Very roughly, this debate is about whether there can be more than one (incompatible) doxastic attitude that can be rendered rational by a single body of evidence. Uniqueness answers in the negative, whereas Permissivism answers in the affirmative. The most plausible form of Permissivism is what has been called ‘interpersonal’ Permissivism. It is the view that it is sometimes rational for distinct individuals to hold different doxastic attitudes in response to the same body of evidence. This contrasts with ‘intrapersonal’ Permissivism, which holds that it is sometimes rational even for a single individual to hold different doxastic attitudes in response to the same body of evidence.⁸

A primary motivation in favour of interpersonal Permissivism is the thought that when two or more people respond differently to the same evidence, they do so because they are relying on different epistemic standards. Since there are multiple rational epistemic standards, then more than one response is rationally permissible (Schoenfield, 2014). As I understand them, epistemic standards are policies that govern what the agent takes evidence to support as well as what and how much evidence is required for the formation or revision of a certain doxastic attitude toward a proposition. Typically, they are not deliberately chosen, nor are they explicitly consulted, and while it may be possible for some agents to bring them into consciousness, it is unusual and cognitively taxing. They are often shaped and varied according to cultural upbringing, and they

⁸ For relevant discussion, see e.g., Kelly, 2013, pp. 303-307 and Wornsip, 2021, pp. 303-08.

might reflect, among other things, different weightings of theoretic virtues (e.g., observational adequacy versus simplicity; Douven, 2009; Willard-Kyle, 2017; Flores, 2021b), different weightings of evidence sources (e.g., ingroup versus outgroup testimony; Balliet et al., 2014; Flores, 2021b), and different epistemic goals (e.g., values placed on avoiding error versus acquiring truths; Fraser, 2020; James, 1979; Kelly, 2013; Flores 2021b). Here are some examples:

- *Car*. On their way out of a restaurant, Jane and Dougie notice that their car is missing. Broken glass is scattered on the ground adjacent to where they had parked it. Jane places a comparatively high value on simplicity relative to other theoretic virtues. Dougie places a comparatively high value on explanatory power. Though they share the same evidence, Jane endorses H1 ('The car was towed away'), which is simple but fits the data poorly, whereas Dougie endorses H2 ('A fleeing bank robber stole the car as a get-away vehicle'), which explains all of the data but it is overly complicated.
- *Campus*. Audrey and Donna are siblings who attend Yale and Harvard, respectively. At a high school reunion, they run into Laura, a common friend and recent Harvard graduate, who brags to Audrey that Harvard has a larger endowment than Yale. Audrey and Donna are more trusting of testimony when it comes from an in-group, rather than out-group, member. While Donna thinks Laura's testimony is very strong evidence of the truth of what she asserts, Audrey finds it weak and untrustworthy.
- *Rumour*. Jeffrey and Sandy hear a rumour that CIA assets blew up the World Trade Centre. Jeffrey has made it a priority to avoid false beliefs about 9/11. He's very cautious about adopting new beliefs on the matter, such that he would only believe the rumour if there was compelling evidence to substantiate it. In contrast, Sandy's main concern is not to miss out on true beliefs about 9/11. She believes the rumour even on scant evidence, for what she fears the most is the prospect of failing to believe it if it turned out to be true.

The individuals in these cases evaluate the same evidence in conflicting ways, depending on the epistemic standards they employ. Jane and Dougie disagree about what their evidence supports because they disagree about the relative weight that should be attached to competing theoretic virtues. Audrey and Donna are differently receptive to Laura's testimony because they differ in whether or not to count Laura as a reliable testifier. Jeffrey and Sandy diverge in their attitudes

about a rumour because they have different levels of tolerance towards epistemic risk: Jeffrey is risk-averse, whereas Sandy is risk-inclined.

A few clarifications are in order to avoid possible misunderstanding of what EM implies. First, an epistemic standard, in the sense I am interested in, has to be something which a thinker *takes* to be reliable at tracking the truth—such that employing it will increase her chances of arriving at truths, or avoiding falsehoods, or both (Shoenfield, 2014). I am not going to undertake the tricky task of spelling out exactly what ‘take’ amounts to. I only want to point out that it need not entail belief. If it did, for a thinker to have an epistemic standard would require her holding a metacognitive belief about its reliability. For example, one would not count as having the epistemic standard ‘If it looks to you as if p , believe p !’ unless one *believed* that believing p when it looks to one as if p makes the truth of p more likely than not. That would rule out the possibility of epistemic standards in creatures like human infants and non-human animals who, arguably, lack the capacity to hold beliefs about their own cognitions. But this seems unduly restrictive. Consider the standard ‘If she walks into the kitchen, believe it is to get you food!’. Without a rationale in hand, nothing prevents us from thinking that a dog who licks its lips and salivates when its owner walks into the kitchen could take such a standard to be truth-conducive, even while being incapable of holding a belief to that effect.

Second, some scholars are content with the notion that we have no justification independent of our standards for choosing them over someone else’s (e.g. Schoenfield, 2014; Willard-Kyle, 2017). Others insist that there must be further standards on which ours rely for their justification (e.g. Feldman, 2006; Goldman, 2010). Fortunately, for present purposes, there is no need to venture into this debate. From the perspective of EM, it does not matter *why* we take our epistemic standards to be truth conducive; it matters only *that* we do. A standard which you take to be an unreliable guide to the truth cannot be of the right sort to count as epistemic. Think, for example, of a standard that recommended against updating upon new information on the sole ground that it would be painful if you did. Other things being equal, however, any standard is preferable if it brings pleasure or avoidance of pain and suffering (or in general some pragmatic advantage).

This brings us to our third point, which is that epistemic standards are to some extent sensitive to context-dependent and practical considerations like how high the stakes are.⁹ For example, in circumstances where grit is required, as when you have invested significant effort in the pursuit of a goal, the best epistemic standard to have will be one which demands much evidence before you update your expectations of success (Morton & Paul, 2019). In the terminology of Morton and

⁹ I say ‘to some extent’, for, as we have just seen, epistemic standards, however practically advantageous they may be, must remain reliable routes to truth by the agent’s lights.

Paul, we may call epistemic standards of this sort the ‘evidential thresholds’ that one employs in a given setting.¹⁰ Variations in the stakes can make a difference in how high one sets evidential thresholds, as well as how one assesses source reliability. An agent might set a high threshold for updating a belief when the costs of acting on a false belief are high (in the example above, giving up personally meaningful goals prematurely), and a lower one when the costs are low. Likewise, if an agent has a great deal at stake, she might be less inclined to rely on another’s testimony as evidence, preferring to find things out for herself. Urgency, too, plays a part here. The need to act promptly might lower evidential thresholds for belief (even when the costs of acting on false beliefs are high) and place more weight on testimony (even without evidence for its reliability).

Fourth, standards may be epistemic without being *actually* useful in arriving at truth.¹¹ We can easily construct scenarios where one thinks a standard likely to lead to truth and the standard is not the least bit truth-conducive. As an extreme case, imagine a person who endorses counter-induction as a reliable mode of inference (Van Cleve, 1984). Such a person has a standard which maps her evidence that all observed *F*s are *G* to a belief that most *F*s are not, or will not be, *G*. So, when told that most students in her class so far have failed the exam, she takes this as evidence that most students in her class will pass the exam. The fact that that this standard is highly unreliable does not detract from its status as epistemic, provided the person thinks it is reliable.

Finally, epistemic standards are not themselves evidence. The reason why one might be prone to think otherwise is that I have not said anything determinate about what evidence is, and so it may seem hasty to dismiss this possibility outright. But regardless of one’s preferred view about evidence, it should be clear that evidence plays a certain functional role, a role that contrasts with the functional role of epistemic standards. At a minimum, a piece of evidence *e* is something which makes a hypothesis *b* more or less probable. Conversely, epistemic standards are criteria whereby to judge, *inter alia*, whether to believe *b* on the basis of *e*. So, both evidence and epistemic standards play a role in determining our attitudes towards some propositions, but they do so in different ways (see, for more on this, Willard-Kyle, 2017).¹²

¹⁰ I will treat evidential thresholds as particular modes of epistemic standards. Morton and Paul (2019) claim that evidential thresholds *ought* to be reliable routes to truth or knowledge, not that they actually are in each and every instance (p. 193). This is compatible with the claim that being an evidential threshold requires that one take it to be reliable at tracking the truth. Do Morton and Paul share my commitment to such a claim? They are not explicit about this point, but it seems very natural to extend their view in this way. At any rate, I adopt this extension here.

¹¹ On the same point, see Ridge, 2018, p. 152

¹² As I was revising the final version of this manuscript, my attention was drawn to Callahan and Titelbaum (forthcoming). Like mine, this paper does some much needed ground-clearing work on the concept of epistemic standards, but it goes into more depth than I could afford to do here. Its main proposal is that agents have their particular epistemic standards by virtue of being *committed* to a system of cognitive rules or procedures. As I understand it, Callahan’s and Titelbaum’s use of the term ‘commitment’ comes close to my use of ‘internalisation’ below (p. 20). But I leave on the agenda of future work the task of specifying the precise extent to which our accounts supplement or complement each other.

4.2. *Details and clarifications*

EM imposes a minimal condition on evidence-responsiveness: the subject must follow an epistemic standard which appropriately explains why she holds the relevant attitude toward p given the evidence available to her. What it means to ‘follow’ a standard is something there is no philosophical consensus on. To delve into this question would take us into deep waters that I cannot even begin to chart here. Three general points seem clear, however. First, one can follow a standard without believing that they do and without holding—or being able to hold (e.g., young kids and non-human animals)—any belief about the standard itself or its content. Suppose Indy, a dog, is weighing evidence bearing on whether some leftover garbage in the street is edible. If Indy is hungry, he will likely use a low evidential threshold. We might then think of him as following a standard that sanctions very weak evidence as sufficient for believing, something of the form:

(Food) If the smell is not foul, believe it is food!

Could Indy and any unsophisticated thinker like Indy even hold a belief as to whether they should believe that p ? It seems unlikely. So, unless we are prepared to deny that unsophisticated thinkers are ever following standards, we should accept that holding a belief about a standard cannot be a necessary condition for following it.

Second, there is an intuitive sense in which someone can accidentally conform to a standard even while failing to follow it (Boghossian, 2008, p. 103). Consider a different dog, Harpo, who believes everything on the ground is food. Like Indy, Harpo will believe that anything which does not smell foul is food, but it would be misleading to say that Harpo is following Food; rather, he just happens to believe in conformity with it. The converse also seems possible. One may follow a standard while failing to conform to it. This would be the case, for example, if one failed to discern whether the conditions specified by the standard are satisfied. Suppose you follow a standard to the effect that you should trust y over z with regard to some relevant evidence, but you mistake y for z , and so you trust z over y . When you form a belief through your trust in z , you are not conforming to your standard, because your belief is based on a source other than that which your standard recommends.

Third, at a first approximation, following a standard is a matter of acting on one's internalisation of that standard.¹³ This can be seen as a way to build a success condition into the notion of following a standard: following a standard is necessarily successful, at least in the sense that if one does not act on one's internalisation of that standard, then one is not really following a standard. EM's clause requiring that there be an appropriate explanatory relation between an epistemic standard x one follows and the attitude one holds is intended to capture and expand on this success condition: one's internalisation of x must be the *reason for which* the attitude is held, which is to say that the explanatory relation between them must be of the right (i.e., non-deviant) sort.

It is extremely difficult to settle on definitions that pin down just what it takes for someone to follow a standard. So far, I have been getting by with the intuitive feel that following a standard involves something more than and different from merely conforming to it. To sharpen our intuitive feel for the distinction, it is helpful to consider a couple of more mundane examples that do not involve dogs or variables. Suppose that some colleagues of Clara are discussing Dylan Farrow's allegations of child molestation at the hand of her adoptive father Woody Allen. And suppose that Clara overhears her boss telling a third party that he believes the allegations are true. Even if she knows little or nothing of the matter, Clara believes the allegations. A possible explanation for her belief is that she follows the epistemic standard:

(Believe Women) If a woman alleges abuse, believe the allegation!

If Clara is to follow Believe Women here, her internalisation of Believe Women must be a reason for which she believes the allegations. If she believed the allegations out of unquestioning respect for her boss, her conforming to Believe Women would be merely accidental and not an instance of following that standard.

When one follows an epistemic standard that one does not conform to, one's internalisation of the standard is the reason for which one believes what one does, even though one's belief is not arrived at by way of conforming to it. To see this, consider Fritz, who follows:

¹³ This seems intuitive enough (Boghossian, 2008, p. 114). If I am following a standard, I must be using it in some sense, and how am I to use it if not by acting on it? Admittedly, the language of 'internalisation' is not much of help with definitory details but it helps capture the intuition that following a standard involves a different relation to it than merely conforming to it. I use 'internalisation' as the most neutral term available to refer to this kind of relation, as the debate about exactly how to understand it is a large one, which I am unable to adjudicate here. See Boghossian (2008, 2014) for a thorough (but sceptical) discussion.

(Health): When it comes to matters of health, prioritise the testimonies of healthcare professionals over those of laypeople!

Suppose that at the same time as receiving Josh's testimony that his ankle pain is from a sprain, Fritz also receives testimony from Jade that it stems from tendonitis. And suppose that he believes that Josh is the physiotherapist, when in reality, it is Jade who fulfils that role, while Josh is simply a shop clerk. Because of this misunderstanding, Fritz trusts Josh's testimony over Jade's, contrary to his own internalisation of Health. Still, it seems true to say that Fritz's internalisation of Health is the reason for which he believes what he does.

I round off this subsection with a clarificatory note. I do not take myself to have provided a full explanation of the requisite condition for EM. I have left the notion of following a standard at a fairly intuitive level, and I have remained neutral about what exactly is involved in a standard's internalisation being explanatory in the right kind of way. So, of course, one will necessarily find many cases where EM is not unambiguously applicable. But the fact that there will be many of these does not undermine my thesis, which is that EM is a minimum threshold that any attitude worth calling evidence-responsive must satisfy. To suppose otherwise is to conflate two different issues: that of establishing minimum requirements for evidence-responsiveness, and that of ascertaining whether they have been met in practice. In the next section, I shall explain my defence of EM.

5 In defence of evidential minimalism

5.1. Low enough

The extension problem revolves around the question of how attitudes based on insufficient evidence and held contrary to evidence can retain enough evidence-responsiveness to qualify as beliefs. What we need to answer this question is an account meeting two desiderata: it must set the bar for evidence-responsiveness (i) low enough to accommodate epistemically defective beliefs but (ii) high enough to demarcate beliefs from other attitudes. I have already argued that RVB and MT have limitations that make them ill-suited (or at least imperfectly suited) to meet these desiderata. I will now argue that EM fares better.

Like RVB and MT, EM makes the suggestion that, though we often believe without evidence and contrary to our evidence, this is compatible with a weakened notion of evidence-responsiveness. The idea is that failures to form beliefs in accordance with the evidence can still be thought of as minimal ways of responding to evidence, provided the relation an attitude entertains with the evidence is regulated by epistemic standards in the relevant sense. On this

understanding, the perseverance of refuted beliefs may be explained in light of the standards that one is following. A standard may recommend maintaining a belief in circumstances where one ought to rationally revise it. (Recall that epistemic standards need not, and all too often do not, reliably produce true beliefs). There are several ways that this can happen.

It could be that the refuted belief was formed on the say-so of a person whose testimony one weighs more heavily than new contrary testimony. It could be that the evidence falsifying the belief is rendered moot by a standard working to undermine the credibility of its source. If the evidence is unfavourable, non-epistemic factors could surreptitiously raise the threshold of evidence needed for belief change or shift standards so that more weight is given to favourable evidence sources. Unreliable inference procedures could have ‘counterevidence’ confirming the ‘refuted’ belief. If one forms beliefs by counter-induction, an enumerative series of F s that are G will be taken as evidence that most F s are not G . These examples are not meant to be exhaustive, but only illustrative of how on EM belief perseverance may often be a principled response to evidence.

RVB and MT are only concerned with how subjects revise existing beliefs or whether they react to detected inconsistencies between them. EM is also concerned with how subjects form beliefs. The fact of being regulated by standards offers a potential rationale for why people often believe without sufficient evidence. For example, Rachel Fraser has recently argued that the epistemic standards reflected in conspiracy thinking are characterised by an attitude of extreme risk-seeking, which she describes as ‘epistemic fear of missing out’ (Fraser, 2020). Someone whose worse fear is losing out on truth will believe even on insufficient evidence, for the possible gain of true beliefs outweighs for them the risk of acquiring false ones. Arguably a person who forms beliefs in this way would be negligent or reckless epistemically. But it would be misleading to say that she is not responding to evidence. It is just that her responses are based on her own standards rather than some objective yardstick.

EM may also bring within the fold of evidence-responsiveness empirical findings that might at first seem in tension with it. I am referring specifically to the illusory truth effect (See above, pp. 1-2). Repetition increases fluency, the experience of ease of processing, and fluency can be a reliable marker for truthfulness (Hassan & Barber, 2021). So, it is not hard to imagine epistemic standards that could recommend treating the experienced ease with which information is processed as evidence for its truth. Nor, accordingly, is it hard to imagine how thresholds could go up for updating beliefs in information that is frequently repeated. I acknowledge, however, that these suggestions are only tentative, all the more so because it remains to be seen whether beliefs formed by familiarity-validity heuristics meet the requisite condition for EM. Indeed, as things

stand, they are perhaps too tentative to translate into an obvious explanatory advantage of EM. Nonetheless, they are suggestions worth making if we are ever to have an account of evidence-responsiveness that applies across the board, to belief formation as well as revision.

A further advantage of EM is that it allows us to avoid some unattractive implications of both RVB and MT. RVB makes some headway against the extension problem, but at the cost of divorcing evidence-responsiveness from how beliefs do in fact behave. All that matters for evidence-responsiveness is that the target attitude be revisable; how it is formed, and why, or why not, it is revised are entirely superfluous. As we have seen in §2 above, this has the unpalatable consequence that an attitude may cease to be evidence-responsive simply by ceasing to be revisable, regardless of the relationship in which it actually stands to the evidence.

EM can avoid being counterintuitive in this way by holding that even unrevisable beliefs may qualify as minimally evidence-responsive if their interactions with evidence are regimented by epistemic standards. Suppose, for example, that one is in a situation in which one maintains a belief despite the availability of evidence to the contrary. Suppose further that one lacks the ability to revise it. On EM, whether the belief enjoys minimal evidence-responsiveness will depend on whether one's not revising it is a direct result of one's incapacity to do so, or whether it can be explained in light of a standard recommending against revision. In the latter case, the person could not have done otherwise, but she maintains the belief because of what her standard recommends, and not because she lacks the capacity to revise it (she would have maintained it even if she could have done otherwise).

To consolidate my view, let me consider one possible objection. Return to my amended version of Sheryl's case, where she is only faced with evidence that overall supports her view that the market is held on Fridays (p. 6). The objection goes as follows. Suppose that at some point after forming her view, Sheryl loses the ability to follow any epistemic standard. Then, as per EM, her view would cease to be evidence-responsive. But if so, this result is just as counterintuitive as RVB's implication that her view would cease to be evidence-responsive upon losing the ability to revise it.

Reply: the fact that one loses the ability to follow epistemic standards need not retroactively affect whether an attitude counts as evidence-responsive. Suppose at t_1 I follow an epistemic standard x which appropriately explains why I form or maintain an attitude A in light of evidence e . EM has it that my attitude would count as minimally evidence-responsive and would thus qualify as a belief. Now say that at t_2 I lose to ability to follow epistemic standards. Relative to e , A would still qualify as evidence-responsive insofar as there is a standard x which I once followed that still appropriately explains A 's relation to e . Of course, were I to obtain new evidence e^* against A ,

then it would no longer be true that \mathcal{A} is evidence-responsive. But—far from being counterintuitive—this is exactly what we want. If there is no standard I follow which can appropriately explain why I retain \mathcal{A} in the face of e^* , then \mathcal{A} 's relation to e^* is far too arbitrary for \mathcal{A} to count as evidence-responsive. The key difference from RVB is that the latter would count \mathcal{A} as unresponsive just by virtue of losing the capacity to revise it, regardless of any changes in evidence.¹⁴

As for MT, we saw that it works with too narrow a conception of minimal evidence-responsiveness. In particular, it has nothing to say about what warrants viewing some cases as minimally evidence-responsive or not. Arguably, there are cases where agents receive countervailing evidence against their beliefs which they ignore and do not take as such. In those cases, there would be no structural inconsistency to react to, for as long as counterevidence is not taken as such, beliefs can retain internal coherence. These cases thus would not violate MT's requirements on minimal evidence-responsiveness. But what is it in virtue of which they are minimally-evidence responsive? After all, there is a clear sense in which they are *not* responding to evidence. An advantage that EM enjoys over MT is that it provides better guidance in this matter. An agent who ignores counterevidence to her belief may still count as minimally evidence-responsive if such behaviour is properly explained by some epistemic standard the agent is following. For example, it has been noted that the epistemic standards of deluded individuals are set so as to overweight the deliverances of one's own senses and underweight the testimony of others (McKay & Mercier, 2023). If true, this would explain why there is no motivation for such individuals to restore coherence (because coherence is never lost), but it would also suggest that by ignoring testimony they are responding to evidence in the way that is best by their own lights.

5.2. *High enough*

Both RVB and MT place a condition on beliefs that is supposed to demarcate them from other attitudes like imaginings and acceptances. Flores claims that imaginings are unlike beliefs in that one can ordinarily fail to revise them in conditions where beliefs would be necessarily revised. The sound of rain tapping on the window will make me revise my belief that it is sunny but need not

¹⁴ An anonymous referee points out that my reply might seem to tie into a suggestion which they find intuitively unattractive, namely, that whether an attitude counts as evidence-responsive constitutively depends on what the environment happens to be like. I must admit that on an intuitive level I do not find this entirely unattractive. (Arguably, whether an attitude counts as knowledge depends on the environment one is in). This is not quite my view, however. I argued that an attitude may cease to be evidence-responsive in the face of a change in the evidence. But that does not mean (and I should not be taken to mean) that a change in evidence itself could determine whether an attitude counts as evidence-responsive. Whether it does depends on there being an epistemic standard which accounts for its behaviour in light of the available evidence.

affect my imagining to the same effect. Of course, we can easily think of cases where imaginings are revised in light of evidence: a child who imagines walking on Mars in a game of make-believe might update her imaginings after seeing NASA's Curiosity Rover's pictures of Martian landscapes. But this is merely contingent matter; it does not derive from the nature of imaginings as such. Bergamaschi-Ganapini instead claims that beliefs exhibit a signature sort of reaction to detected inconsistencies, which is alleged to mark them apart from other attitudes. As we have seen, though, this criterion would at best enable us to tell beliefs from imaginings, but not beliefs from acceptances.

Structurally, my proposal for a demarcation criterion is parallel to that of Flores. I point to a feature that beliefs share with other attitudes, but which holds only contingently of the latter. I claim that imaginings and acceptances are unlike beliefs in that, although they can be regulated by epistemic standards, this is contingent upon the particular goal they are being used for, and the reasons for engaging in the pursuit of that goal.

Start with imaginings. There are different uses to which imagination can be put that are subject to constraints (Kind & Kung, 2016; Noordhof, 2018). For example, if my aim is to imagine something pleasing for myself, I will model my imagining after what I take to be pleasing. In that case, when I engage in an episode of pretence that is not pleasing enough by my lights, the constraint in question may recommend shifting the pretence to something more pleasing. Many other examples could be given, but the general point is clear; we not only initiate imaginings toward specific goals, we also constrain our imaginings to fit those goals.

If the goal sought after is some non-epistemic good like pleasure, the constraints operating on imagining will be non-epistemic, too. Notice that I am not denying that what one imagines here might be informed by what one knows about the actual world (Van Leeuwen, 2023, p. 47). As I fantasise about, say, petting a wolf, I might recreate, in my imagination, past experiences of seeing real wolves on TV. But whether I stick closely to reality or not is not dictated by the goal of imagining something pleasant; it is incidental to it. I am free to imagine whatever I like (realistically or unrealistically) so long as it pleases me.

Matters are different when imaginings are employed to epistemic ends, such as the fixation of counterfactual beliefs (cf. Kind, 2016). One example is when we try to imagine what might happen if we take a particular course of action. Say that I am trying to decide whether to drive down a snowy road, and to that end, I try to imagine what might actually happen if I do. This is a plausible sort of case in which the nature of my goal dictates that I imagine realistically. In the present case, that would require, among other things, that I imagine my car with the attributes it

actually has. So, for instance, if I search and cannot find snow chains in the trunk, I should adjust my imagining to incorporate that evidence.

The standard in question seems epistemic. I cannot learn whether my car will handle snow if I imagine a car different from mine. I expect that if I imagine my car as it appears in all relevant respects, then this will increase the chances that my imagining represents the future correctly. The question arises: if epistemic standards apply to imaginings too, why think they are a ‘mark’ by which beliefs can be distinguished from imaginings? The answer is that there is a difference in the way beliefs and imaginings are subject to epistemic standards.

Imaginings are subject to epistemic standards when we set out to use them in pursuit of an epistemic goal, like figuring out what is likely to be the case in some circumstances. Every imagining is such that it is generally possible to control it in the service of intended goals. In this way, we can indirectly determine the standards through which these goals can be attained. Beliefs are different, however. Whatever their epistemic goal, it is one over which we have no deliberate control, and not the sort of goal which we might or might not have. This explains why the standards governing beliefs do not vary according to which goals are contingently willed; they remain stable over time.

Similar points can be made about acceptances. When I accept p as a premise for reasoning in a context, I am acting intentionally to achieve a desired goal. Like with imaginings, goals can differ. In the bookstore example, I remain unconcerned with getting at the truth. Learning that the bookstore is closed would get me to stop accepting that it is open. But this does not involve any kind of truth-aiming. I relinquish my acceptance because it is no longer useful to me, not because it is false. We can think of cases, though, where one’s goal is truth. For instance, in the context of figuring out whether p , a scientist might accept p in order make predictions about things she would expect to observe if p and design experiments to test these predictions. Given her current goal, I take it as fairly clear that she would be rationally required to reject p in the face of compelling evidence that p is false. Should not we then say that her acceptance is subject to an epistemic standard of some sort? Yes, but here too there is reason to be careful in drawing general conclusions. When acceptances are so governed, this is a merely contingent fact about them (i.e., it only happens when we intend to use them for the sake of truth), and thus something importantly different from what we see in beliefs.

6 Conclusion

Most existing theories on the nature of belief are committed to the claim that beliefs generally and characteristically respond to evidence. However, many (apparent) beliefs are not so responsive, at

least not in the sense that philosophers have tended to suppose. This raises the question of what the manner and extent of responsiveness must be in order for an attitude to qualify as belief. If we leave room for degrees of responsiveness, we might be able to identify a minimal degree of responsiveness that is high enough to mark belief off from other attitudes, yet low enough to admit of imperfections in the relation between belief and evidence. I have argued for such a basic minimum to be defined in relation to epistemic standards. An attitude A counts as minimally evidence responsive if, and only if, the agent who has it follows an epistemic standard which properly explains A 's relation to the evidence. The resulting view allows that attitudes may retain responsiveness in this minimal sense even when we ignore evidence we have against them, or when we are incapable of revising them. At the same time, the view provides a principled way to distinguish beliefs from other related cognitive attitudes. Even though beliefs are not the only attitudes whose evidence relations are subject to regulations via epistemic standards, beliefs and non-beliefs are subject to distinct modes of regulation.¹⁵

References

- Adler, J. E. (2002). *Belief's own ethics*. Cambridge, MA: MIT Press.
- Anderson, C.A., Lepper, M. R. & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology*, 39, 1037–1049. doi: doi.org/10.1037/h0077720.
- Balliet D., Wu J., & De Dreu C. K. (2014). *Ingroup favoritism in cooperation: a meta-analysis*. *Psychological Bulletin*, 140, 1556-81. doi: 10.1037/a0037737.
- Begby, E. (2021). Evidential Preemption, *Philosophy and Phenomenological Research*, 102, 515-530. doi: 10.1111/phpr.12654
- Begby, E. (2024). From a Belief Polarization to Echo Chambers: A Rationalizing Perspective, *Episteme*, 21, 519-539. doi: 10.1017/epi.2022.14.

¹⁵ I owe my greatest debt of gratitude to Matt Parrott for detailed feedback and discussion at every stage of writing this paper. Additional thanks go to Lisa Bortolotti and Carlo Cordasco for comments on past drafts, and to an anonymous PPR referee for their useful suggestions. I have also benefitted from conversations with Bogdan Dicher, Carolina Flores, David Horst, Jessica Moss, and Joshua Rowan Thorpe. Work on this paper was supported by the Fundação para a Ciência e a Tecnologia (FCT) — grant number 2023.07700.CEECIND.

Bendana, J. & Mandelbaum, E. (2021). The Fragmentation of Belief. In D. Kinderman, A. Onofri & C. Borgoni (Eds.), *The Fragmented Mind* (pp. 78-107). Oxford: Oxford University Press.

Bergamaschi-Ganapini M. (2020). Belief's minimal rationality, *Philosophical Studies*, 177, 3263-3282. doi: 10.1007/s11098-019-01369-y.

Bergamaschi-Ganapini M. (Forthcoming). A Minimalist Threshold for Epistemically Irrational Beliefs, In E. Schwitzgebel & J. Jong (Eds.), *The Nature of Belief*. New York: Oxford University Press.

Boghossian, P. A. (2008). Epistemic rules, *The Journal of Philosophy*, 105, 472-500. Reprinted in P. A. Boghossian (2008). *Content and Justification* (96-120). Oxford: Oxford University Press.

Boghossian, P. A. (2014). What is inference? *Philosophical Studies*, 169, 1-18. doi:10.1007/s11098-012-9903-x.

Bortolotti, L. (2009). *Delusions and Other Irrational Beliefs*. Oxford: Oxford University Press.

Bratman, M. E. (1992). Practical reasoning and acceptance in a context, *Mind*, 101, 1-15. doi: 10.1093/mind/xcviii.391.367.

Callahan, L. & Titelbaum, M. (Forthcoming). What are epistemic standards? In S. Goldberg & M. Walkers (Eds.), *Attitude in Philosophy*. Oxford: Oxford University Press.

Chan, M. S., C. R. Jones, K. H. Jamieson, & D. Albarracín. (2017). Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychological Science*, 28, 1531–1546. doi: 10.1177/0956797617714579.

Cohen, L. J. (1989). Belief and Acceptance, *Mind*, 98, 367-89. doi:10.1093/mind/xcviii.391.367.

Currie, G. & Ravenscroft, I. (2002). *Recreative Minds: Imagination in Philosophy and Psychology*. Oxford: Oxford University Press.

- Dorst, K. (2023). Rational polarization, *Philosophical Review*, 132, 355-458. doi: 10.1215/00318108-10469499.
- Douven, I. (2009). Uniqueness revisited, *American Philosophical Quarterly*, 46, 347–361.
- Egan, A. (2008). Imagination, Delusion, and Self-deception. In T. Bayne & J. Fernandez (Eds.), *Delusion and Self-deception: Affective and Motivational Influences on Belief Formation*. Hove: Psychology Press.
- Fazio, L. K., Rand, D. G., & Pennycook, G. (2019). Repetition Increases Perceived Truth Equally for Plausible and Implausible Statements, *Psychonomic Bulletin and Review*, 26, 1705–1710. doi: 10.3758/s13423-019-01651-4.
- Feldman, R. (2006). Epistemological puzzles about disagreement, In S. Hetherington (Ed.), *Epistemology Futures* (216-236). Oxford: Oxford University Press.
- Feldman, R. (2010). Reasonable religious disagreements, In L. Anthony (Ed.), *Philosophers without Gods: Meditation on atheism and secular life* (194-214) New York: Oxford University Press.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Flores, C. (2021a). Delusional evidence-responsiveness, *Synthese*, 199, 6299-6330. doi: 10.1007/s11229-021-03070-2.
- Flores, C. (2021b). Epistemic Styles, *Philosophical Topics*, 49, 35-55. doi: 10.5840/philtopics2021149214.
- Flores, C. (forthcoming) Resistant beliefs, responsive believers, *The journal of philosophy*.
- Fraser, R. (2020). Epistemic FOMO, *The Cambridge Humanities Review*, 16, 15-29.
- Fricker, M. (2007). *Epistemic injustice: power and the ethics of knowing*. Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780198237907.001.0001.

Gendler, T. S. (2008). Alief in action (and reaction). *Mind & Language*, 23, 552-585. doi: 10.1111/j.1468-0017.2008.00352.x.

Glüer, K., & Wikforss, Å. (2013). Aiming at truth: On the role of belief, *Teorema*, 32, 137-162.

Goldman, A. (2010). Epistemic relativism and reasonable disagreement. In R. Feldman & T. A. Warfield (Eds.), *Disagreement* (188-215). Oxford: Oxford University Press.

Harmon-Jones, E. & Mills, J. (2019). An introduction to cognitive dissonance theory and an overview of current perspectives on the theory. In E. Harmon-Jones (Ed.), *Cognitive dissonance: Reexamining a pivotal theory in psychology* (3-24). Washington DC: American Psychological Association.

Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 16, 107–112.

Hassan, A., & Barber, S. J. (2021). The effects of repetition frequency on the illusory truth effect. *Cognitive Research: Principles and Implications*, 6, 1-12. doi: 10.1186/s41235-021-00301-5.

Helton, G. (2020), If you can't change what you believe, you don't believe it, *Nous*, 54, 501-526. doi: 10.1111/nous.12265.

James, W. (1979). *The will to believe and other essays in popular philosophy*. Cambridge, MA: Harvard University Press.

Jelalian, E. & A. G. Miller. (1984). The perseverance of beliefs: Conceptual perspectives and research developments. *Journal of Social and Clinical Psychology*, 2, 25–56. doi: 10.1521/jscp.1984.2.1.25.

Kelly, T. (2013). Evidence can be permissive In In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (298-311). New York: Blackwell.

Kind, A. (2016). Imagining under constraints. In A. Kind & P. Kung (Eds.), *Knowledge through imagination* (145-159). Oxford: Oxford University Press.

Kind, A. & Kung, P. (2016). Introduction: The puzzle of imaginative use. In A. Kind & P. Kung (Eds.), *Knowledge through imagination* (1-37). Oxford: Oxford University Press.

Lacassagne, D., Béna, J., & Corneille, O. (2022). Is Earth a perfect square? Repetition increases the perceived truth of highly implausible statements, *Cognition*, 223. doi: 10.1016/j.cognition.2022.105052.

Levy, N. (2015). Neither Fish Nor Fowl: Implicit Attitudes as Patchy Endorsements, *Nous*, 49, 800–823. doi: 10.1111/nous.12074.

Levy, N. (2022). *Bad Beliefs: Why They Happen to Good People*. Oxford: Oxford University Press.

Mandelbaum, E. (2019). Troubles with Bayesianism: An introduction to the psychological immune system. *Mind & Language*, 34, 141–157. doi: 10.1111/mila.12205.

McKay, R., & Mercier, H. (2023). Delusions as Epistemic Hypervigilance, *Current Directions in Psychological Science*, 32, 125-130. doi: 10.1177/09637214221128320.

Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, Mass: MIT press.

Neta, R. (2019). The basing relation, *Philosophical Review*, 128, 179-217. doi: 10.1215/00318108-7374945.

Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, 17, 141-161. doi: 10.1017/epi.2018.32.

Noordhof, P. (2018). Imaginative content, In F. Macpherson & F. Dorsch (Eds.), *Perceptual Imagination and Perceptual Memory* (96-129). Oxford: Oxford University Press.

Paul, P. & Morton, J. (2019). Grit, *Ethics*, 129, 175–203. doi: 10.1086/700029.

Ridge, M. (2018). How to be an epistemic expressivist. In C. McHugh, J. Way, & D. Whiting (Eds.), *Metaepistemology* (141-158). Oxford: Oxford University Press.

- Ross, L., Lepper, M. R. & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology*, 32, 880–892. doi: 10.1037/0022-3514.32.5.880.
- Salice, A., & Miyazono, K. (2020). Social epistemological conception of delusion, *Synthese*, 199, 1831-1851. doi: 10.1007/s11229-020-02863-1.
- Schoenfield, M. (2014). Permission to believe. Why permissivism is true and what it tells us about irrelevant influences on belief, *Nous*, 48, 193-218. doi: 10.1111/nous.1206.
- Shah, N, & Velleman, J. D. (2005). Doxastic Deliberation. *The Philosophical Review*, 114 , 497–534. doi: 10.1215/00318108-114-4-497.
- Van Cleve, J. (1984). Reliability, Justification, and the Problem of Induction, *Midwest Studies in Philosophy*, 9, 555-567. doi: 10.1111/j.1475-4975.1984.tb00077.x.
- Van Leeuwen, N. (2014). Religious credence is not factual belief. *Cognition*, 133, 698-715. doi: 10.1016/j.cognition.2014.08.015.
- Van Leeuwen, N. (2023). *Religion as make-believe. A theory of belief, imagination, and group identity*. Cambridge, MA: Harvard University Press.
- Velleman, J. D. (2000). On the Aim of Belief. In J. D. Velleman, *The Possibility of Practical Reason* (123–143). *New York: Oxford University Press*.
- White, R. (2005). Epistemic permissiveness, *Philosophical Perspectives*, 19, 45-459. doi: 10.1111/j.1520-8583.2005.00069.x.
- Wedgwood, R. (2023). *Rationality and Belief*. Oxford: Oxford University Press.
- Westfall, M. (2024). Polarization is epistemically innocuous, *Synthese*, 204, 1-22. doi: 10.1007/s11229-024-04739-0

Willard-Kyle, C. (2017). Do great minds really think alike? *Synthese*, 194, 989-1026. doi: 10.1007/s11229-015-0984-x.

Worsnip, A. (2021). *Fitting things together. Coherence and demands of structural rationality*. New York: Oxford University Press.