

In Defence of Ontic Austerity for Belief

The contributions to this volume all give answers to the question *what is belief?* Many extant and developing accounts of belief hear the question as one of modal strength – that is, as picking out the nature of belief across worlds – and theorise in such terms. I argue that if we want our account of belief to be explanatorily adequate, we should separate the modal question concerning the very nature of belief, from a more local question concerning the manifestation of its various guises in the actual world. I propose that what is necessary to belief is its motivational role, but that this is not where our explanatory gains lie. Rather, the contingent biological circumstances of belief’s manifestation in the actual world is where we find explanations of the behaviour of garden variety beliefs, as well as the resources to accommodate a couple of *edge cases* (i.e. religious belief and self-deceptive belief). I close by identifying three advantages delivered by my ontically austere account: unification, explanatory power, and parsimony.

1. The belief theorist’s dilemma

The job of giving an account of belief is an unenviable one. In the most basic version of the task, we want to identify a property or set of properties x . And we want x to allow us to do two things: *unify* and *explain*. That is, we might have pre-theoretical intuitions about which attitudes are deserving of doxastic status, and we want x to do the work of collecting all and only those together under the heading of *belief*, such that every member of the set shares the defining property of that set. We want x to be unifying. But unification alone doesn’t satisfy us. We also demand of x that it explains the various ways members of the set behave, most notably, their conditions of manifestation and extinction, and their overall role in our cognitive architecture. That is, we want x to be the answer to questions like *why can’t we believe at will? Why does undeniable counterevidence extinguish belief?* We want x to be explanatory.¹

It can seem, however, that we must pick: unify *or* explain? If we pick *unify*, we must be prepared to go cheap. As Neil Van Leeuwen and Tania Lombrozo have pointed out:

Pick a property to help define ‘belief,’ and a motley crew of delusions, ideological beliefs, religious beliefs, and so forth, wrecks havoc on your definition. (2023: 2)

In pursuit of unification one must be in the market for a feature which captures unremarkable garden variety attitudes regarding train times and weather reports, as well as attitudes regarding deities and alien abduction. But in stripping belief to its bare bones to tolerate the range of its tokens, we will be hard pressed to put that feature to much explanatory work.

¹ It might be wondered why we should expect the work of delineation and explanation to be done by the same property, especially since this doesn’t seem to match any expectations about other epistemological notions (e.g. it wouldn’t be standard to think that capturing the nature of knowledge by appeal to some x thus gives us the resources to explain its functioning). I’m taking my cue here from folk who seek to explain various features of belief (those I overview in §2), and do so by appeal to the very nature of belief. Normative and teleological account exemplify this approach, where the identified norm or aim claimed to be constitutive of belief is also what is taken to explain the features of interest.

Van Leeuwen and Lombrozo advocate moving away from the idea that belief is one thing, and instead being open to our pre-theoretical use of the term *belief* capturing processing involving distinct cognitive attitudes (2023: 3). The lack of consensus surrounding the answer to this volume's question is diagnosed as arising from the 'misplaced expectation' that there is a single psychological kind responsible for all instances of believing. We should stop the search for a single unifying kind properly captured as belief, and instead recognize that:

a multi-dimensional property space within which different varieties of believing – along with other cognitive attitudes – can be characterized, and the interesting natural kinds will be the clusters (and clusters of clusters) that emerge in that property space through empirical investigation. (2023: 3)

Once we give up on unification, we can of course characterize varieties of believing much more robustly, adding features explanatory of that particular variety's behaviour.

In this chapter I argue that we can have our cake and eat it too. Van Leeuwen and Lombrozo's lesson is well taken, and my position is friendly to it. When asking what the relevant clusters might look like, they suggest that 'they will align with different functions or aims that the beliefs serve' (2023: 3). My view captures beliefs as produced in the service of different biological functions. We will see that characterizing beliefs as having various biological etiologies allows for great explanatory power. However, let us not give up on unification! We can instead opt for ontic austerity in capturing what unites all of belief's instances across worlds. Now, admittedly, in satisfying the desire for unification, we do not satisfy the explanatory goals that such unification has been thought to promise (§4). There is, I say, no x that will do that. The victory of unification, then, might be thought Pyrrhic. However, all is not lost. We should leave the messy business of everyday variation and the theory-pressuring motley crew of edge cases to be picked up by the diversity of (contingent) ways in which token instances of belief manifest in the actual world. That is, when answering the question *what is belief?* the philosophical goals of *unification* and *explanation* do not arise from the same theoretical resources, but can nevertheless both be served by one overall approach. We should attend to the contingent biological circumstances in which beliefs are formed, but ensure that we resist taking features of belief arising from such circumstances to be necessary ones. Explanations for belief's behaviour can be found in its manifestation in the actual world, set against a background motivational account.

2. The belief-truth link

Having a clear conception of belief might allow us to better understand the constraints under which it operates. Some such constraints are often taken to follow from belief's relationship to truth, and so one way to approach our task is by getting clear on what that relationship looks like. As Bernard Mayo put it, '[c]entral to any discussion of belief is the question of the relation between belief and truth' (1964: 139). This, then, will be my approach. In this section I identify three features of belief that point to such a relationship. Later I show that my account can explain each of them (§5.2).

The first is transparency in doxastic deliberation. That is, ‘when asking oneself *whether to believe that p*’ one must ‘immediately recognize that this question is settled by, and only by, answering the question *whether p is true*’ (Shah 2003: 447). Transparency is not merely that the truth of *p*, by the subject’s lights, is required for her to believe that *p*. It is rather that recognizing the truth of *p* in deliberation is sufficient for believing it.² Of course, transparency doesn’t rule out belief being influenced by factors unrelated to truth, it is only that ‘one cannot deliberately, and in full awareness, let one’s beliefs be guided by anything but truth’ (Steglich-Petersen 2006: 503).

The second feature is our inability to bring about beliefs at will. That is, ‘I cannot bring it about, just like that, that I believe something’ (Williams 1973: 148). We can be a little more careful and follow Paul Noordhof’s characterization of what it is we cannot do: ‘unmediated conscious belief-production is impossible’ (2001: 248). We cannot, as a result of a conscious intention to believe that *p* now, go ahead and do so without a mediating act or evidence, where that conscious intention and the resulting belief stand in the same relation to each other as do intention and action in the context of intentional actions (2001: 248). The final feature is belief’s standard of correctness. Truth and falsehood are a ‘dimension of assessment of beliefs as opposed to many other psychological states or dispositions’ (Williams 1973: 136). It is typically thought that the standard of correctness for belief is that *a belief is correct if and only if it is true*.

Some approaches to explaining the above features appeal to belief as constitutively normative, and understand the claims of transparency and doxastic involuntarism to be normative claims (e.g. Shah 2003, Shah and Velleman 2005), and the standard of belief to be a normative standard. Understanding the features in this way is not common ground (see e.g. Steglich-Peterson 2006), nor is it to be taken for granted in setting out our explanatory task. With respect to the first two features, I will take the explananda to be descriptive ones, and with respect to the standard of correctness, I will take this to be a non-normative standard. This latter move may strike many as unsatisfying, since it misses out on the standard’s presumed normative oomph. I note then that standards need not be normative ones. As Krister Bykvist and Anandi Hattiangadi have argued:

Judging that ϕ -ing is correct is compatible with judging that one ought not to ϕ . Judging that ϕ -ing is incorrect is compatible with judging that one ought to ϕ . When it is a fact that ϕ -ing meets a certain standard, there is always a further question whether the standard ought to be met. In some cases, the standard ought to be met, in others, not. (2013: 103)

Standards alone, then, do not generate obligations without an additional element which makes those standards normative ones. Elsewhere I have argued that the normative-reading of this standard has little explanatory utility (see Sullivan-Bissett 2018, Sullivan-Bissett and Noordhof 2020). I have also given an error account of the perceived epistemic normativity thought to characterise this standard (Sullivan-Bissett 2017a), and have argued that we have no duties or obligations to believe

² Some philosophers deny that transparency characterises our doxastic deliberation (see e.g. McHugh 2012, 2013, McCormick 2015). I have responded to these arguments elsewhere (2017b, 2017c).

(Sullivan-Bissett 2020). I will not rehearse these arguments but will take the explanatory burden to be discharged if my account can explain belief's standard of correctness in this thinner, non-normative sense.

In what follows I show that my ontically austere approach has the tools to explain the relationship between belief and truth indicated by these three features, before moving on to edge cases. I begin by speaking to the essential nature of belief.

3. What is belief? The modal background

Although providing unification, the modal question regarding belief is the less interesting and the less explanatorily fruitful one. On my view (drawing from a familiar framework, see e.g. Stalnaker 1984, Davidson 1963), belief is an attitude with a particular motivational role. For our purposes we can follow a characterization offered (but not endorsed) by Lucy O'Brien: belief is an attitude which, by itself, and relative to one's background desires, disposes the subject to behave in ways which would fulfill her desires, if the contents of the belief were true, across contexts (O'Brien 2005: 56). The 'across contexts' is my addition, for reasons that will soon become clear. If you're wondering what you'll run into when you come across beliefs in modal space, you'll find a state which is *motivational*. If you want to capture belief across worlds: *action is where the action is*. Beliefs are, 'maps of the world in the light of which we are prepared to act' (Armstrong 1973: 4), in pursuit of desire satisfaction, across contexts.

An objection might immediately occur to the reader: characterizing belief as a state which motivates action will not distinguish it from, for example, imagination. David Velleman argues that belief's motivational role is not unique to it. His example is that of a child imagining that *she is an elephant*, and he claims of her that she will be disposed to act in ways such that, if the content of the imagining were true, it would promote the satisfaction of her desires (2000: 255–6). This state of imagination, then, will equally meet any description of belief by appeal to motivational role alone (see also Ichino 2019). If that's right, then our modal condition won't delineate belief after all.

However, this objection does not work. Many authors have noted some version of the idea that the role that imagination plays in motivating action is context dependent, and it is because I recognise that some non-belief states have motivational power that I include *across contexts* in my characterization of the motivational role definitive of belief. The roles played by imagination and belief will come apart in some contexts. Enumerating those allows us to distinguish belief from imagination (see Noordhof 2001: 252, O'Brien 2005: 59, and Glüer and Wikforss 2013: 143–5).³

³ Some folk have resisted the idea that we can neatly distinguish belief from imagination. Susanna Schellenberg (2013) has appealed to imaginative immersion as involving seamlessly transitioning from an imagining to something more belief-like, to argue that belief and imagining are on a continuum. Andy Egan (2008) has appealed to delusions as a case which neither matches the functional profile of belief nor imagining, but instead partly matches each. He takes it that this clears space for the continuum hypothesis. If I am right that we can delineate belief in an ontically austere way, the continuum hypothesis might be harder to motivate (given there are fewer constitutive features that a candidate state could exemplify to a belief-like degree). For the purposes of this paper I will proceed on the assumption that the case for the continuum hypothesis has not been made (for excellent arguments for this claim see Kind 2023).

For readers already penning other counterexamples to my characterization of belief's motivational role, I am not the sparring partner you're looking for. I will not spend time here seeking to precisify the exact motivational role of belief, or even further defending the one given, since what I want to say is that however it is captured, that's the extent of what we should say about the necessary nature of belief. Belief is, at base, an attitude that makes us *do stuff*, in concert with other mental states (notably, desires), *across contexts*.⁴ Of course, a suitable *ceteris paribus* clause should also be read in. We proceed then with the modal question answered in an ontically austere way.

4. Limitations of the modal view

I do not see much promise in accounts of belief which seek explanatory gains via some unifying modal feature of belief (see fn. 5), but it is worth quickly outlining why that is definitely a non-starter when it comes to my preferred view.

Let us return to transparency, which will be enough to see the explanatory limitations of the modal claim. Can an account of belief which captures it as a state with a particular motivational role explain this feature of doxastic deliberation? Might there be any mileage in the idea that given belief's motivational role, true beliefs are going to be more useful than false beliefs? Jerome Dokic and Pascal Engel in their discussion of Ramsey make the point:

a false belief can only be useful locally and by accident. A false belief cannot properly guide our behaviour in every possible situation. In contrast, truth systematically promotes the success of action, on a large scale. Truth offers a *guarantee* of success that falsity cannot hope to deliver. (2002: 48)

So we might say something like this: our deliberative belief formation being characterized by transparency is good news for believers. If deliberation were not transparent to considerations of truth, I might end up with more false beliefs and, in turn, engage in fewer or less successful or actions. So transparency is present because I know that having true beliefs is more likely to bring about successful action.

This explanation is no good. We will see shortly that the relationship between belief's motivational role and successful action should indeed inform our overall account (§5.1). But in framing transparency as an agential phenomenon, our present explanation implies that I have control over the considerations that my deliberative belief formation is sensitive to. But if this were true, I would be able to answer the question *whether p is true* in the positive, without *thereby* believing that *p* (if I did not

⁴ Perhaps it will be granted that the particulars of the downstream role need not be discussed here, given the theoretical role of the motivational account in my overall view, but the profile of belief understood in a downstream way is a decision point that calls for justification. Why characterise the role of belief in terms of its downstream causal role in action, rather than its upstream causal role, e.g. being produced by perception, or its internal causal role, e.g. in inference? For now I note only that I am inclined to think that an upstream causal role for belief will not do the job of delineation. For example, *being produced by perception* will not be true of all beliefs, and this feature may well be shared by some cases of imagination (where I am prompted to imagine by my surroundings). Similarly, imaginative states might play a causal role in inference (see Sinhababu 2013: 160–1), and there is some evidence that implicit biases partake in inferential relations (see Mandelbaum 2016).

care for desire satisfaction, for instance). The explanation gives us a gap between *whether p is true* and *whether to believe that p*. But there is no such gap in doxastic deliberation. The claim that belief is essentially motivational does not have the resources to explain why this is the case.⁵

So is all doomed for the project of resisting ontic bulge for belief? I think not. Although I am strongly in favour of keeping the essential nature of belief ontically thin, I have no such qualms when it comes to capturing the contingent nature of belief as it presents in the actual world. It is here that our account of belief can be genuinely illuminating.

5. What is belief? The contingencies

I propose a programme which pays attention to the biological circumstances in which beliefs are formed, but which resists any ontological reading of belief's features which arise from those circumstances, lest we mistakenly project the particular circumstances of our biological heritage into a modal claim. It is belief's particular manifestation in the actual world, understood against a background of belief as essentially motivational, where we will find explanations for various ways belief and belief formation work, both in central cases, and at the edges.

5.1 Function one

On my view, it is one of the biological functions of our mechanisms for belief production which explain transparency, our inability to believe at will, and belief's standard of correctness. This is the function to *produce true beliefs*. Functions here should be understood historically, that is, as picking out something performed by a trait's ancestors which 'helped account for the proliferation of the genes responsible for it' and thus 'helped account for its own existence' (Millikan 1989: 289). This function ascription to our mechanisms of belief production is, of course, neither new nor surprising (see similar claims in Papineau 1987, Millikan 1993). These mechanisms have proliferated because—given the essential nature of belief (its motivational role)—creatures with true beliefs enjoyed biological advantages which increased inclusive fitness: they were more likely to survive because their actions were successful.

Ascriptions of biological functions do not amount to reliable generalizations or even statistical likelihoods, but in cases where functions are not performed, questions are licensed. It is useful at this point to introduce some further theoretical apparatus. Ruth Millikan's sense of *normative historical normalcy*, signified with capitalization, is distinct from *statistical normalcy*. To demonstrate: sperm Normally fertilize ova, but it is not the case that sperm normally fertilize ova (Millikan 1984: 34). In cases of non-

⁵ What about the prospects for accounts of belief which *build in* a necessary relationship to truth? For example, teleological aim-based accounts might have it that belief aims at truth, and given this, deliberation over whether to believe that *p* is structured by the aim one has in forming a belief. Alternatively, normative accounts of belief might have it that in forming a belief, our activity is governed by a norm, something like *believe that p only if p is true*. Elsewhere I have argued that both kinds of explanation also fail (Sullivan-Bissett 2018, Sullivan-Bissett and Noordhof 2020). In brief, both accounts make the same mistake: in casting transparency as an agential phenomenon, they leave a gap where there is none.

performance of a given biological function, we might say that the trait token in question *malfunctioned*, or we might say that the circumstances for functional performance were abnormal, and the trait *misfunctioned*. I will say more about the difference later (§6.1).

Finally, let us turn to a more specific kind of function attribution: *relational proper function*. When we say of a device that it has a function of this kind, we are saying of it that ‘its function [is to] produce something that bears a specific relation to something else’ (Millikan 1984: 39). Relational proper functions are the effects that have helped account for the selection of the producing mechanisms (Millikan 1984: 26). That produced by these mechanisms are *adapted devices*, and that to which the adapted device is adapted is the *adaptor* (Millikan 1984: 40).

We can now formulate our function ascription more precisely: our mechanisms of belief formation have the relational proper function of producing true beliefs. Such mechanisms have proliferated because—given the essential motivational role of belief—creatures with true beliefs will be more likely to survive because their actions will be successful. The modal character of belief then does *some* explanatory lifting in its providing an explanation for why, in the actual world, mechanisms directed at truth are biologically adaptive (because true beliefs tend to facilitate successful action). In the actual world, beliefs produced by these mechanisms are *adapted devices*, and the mechanisms which produce them Normally produce true beliefs.

5.2 Explaining the belief-truth link

Let us return to the three features outlined earlier (§2). When a subject poses to herself the question *whether to believe that p*, she attends to the question *whether p is true*. In functional terms: transparency is achieved by certain causal facts which hold for our mechanisms of belief formation. These causal facts make it such that whenever these mechanisms perform their relational proper function, they are sensitive to the way the environment is, that is, the adaptor for the produced belief. This is understood at the agent level as epistemic considerations pertaining to the truth of *p*. Instances of epistemic considerations settling whether to believe are instances of believers immediately and inescapably attending to the environment to which the belief is adapted. The move from one question to another here is not achieved by anything the agent *does* (as in other accounts, see fn. 5). Why do these causal facts which secure transparency obtain? Why is our cognitive architecture thus arranged? Here we can appeal to the evolutionary history of our mechanisms for belief production: the cognitive architecture which secures transparency has been selected for its role in producing true beliefs (for more details see Sullivan-Bissett 2018).⁶ Transparency thus

⁶ Does this explanation make my account a pragmatist one? The background notion of belief operative in one’s explanation of transparency makes a difference to which accounts of epistemic reasons such an explanation is consistent with. An evidentialist will of course assign the explanatory burden of transparency to their preferred theory of epistemic reasons, set against a background commitment to belief being essentially connected to truth, via a constitutive aim or norm. Things are straightforward for the pragmatist, since her view regarding reasons for belief does not so easily lend itself to explaining transparency (that explanation is to be found elsewhere). (see Dular and Fortier 2022 who offer an evolutionary account of transparency so as to show that transparency does not help us adjudicate between evidentialism and pragmatism). My account of belief is not one which generates norms of good belief regulation. As such, it is silent on the debate between evidentialists and pragmatists.

comes out as contingently characterizing deliberative belief formation. My explanation leaves open the metaphysical possibility of beliefs arising deliberately but non-transparently.⁷

Next is our inability to believe at will, captured earlier by Noordhof's *uncontrollability thesis*, that is, 'unmediated conscious belief-production is impossible' (2001: 248). The truth of this thesis falls straightforwardly out of doxastic deliberation being governed by transparency. If beliefs are adapted to the environment, then to produce a belief by an effort of the will alone would be to produce a belief irrespective of the environment to which it is adapted. The causal facts which secure transparency will rule out beliefs being formed in this way.

Finally, I accommodate belief's standard of correctness by noting that the function proper to our mechanisms of belief production provides the only sense in which true beliefs are correct and false beliefs are incorrect. The contents of belief can vary with respect to their meeting the correctness conditions laid down by biology (for more on how my function ascriptions interact with the standard of correctness for belief, see my 2017a, §3, §4). I refer readers dissatisfied with this explanation to my above remarks concerning the possibility of non-normative standards (§2).

In the preceding, I have (albeit briefly) sought to show that my account can explain the features indicative of a relationship between belief and truth. I turn now to the second function claim of my account.

5.3 Function two

True belief production is not the only thing at which our mechanisms for belief-production are aimed. There are lots of cases of false belief which seem to have been produced by mechanisms doing what they are *supposed to do*. That is, producing adaptive beliefs which are usually false. If there are such cases, we have good reason to specify a second function proper to those mechanisms, so we can account for the etiology of these cases

Given this, I suggest that the second function proper to our mechanisms of belief formation: the production of *useful* beliefs. Usefulness here should be understood not as it relates to approximating to truth, but rather useful insofar as they facilitate self-organization, maintain self-esteem, avoid psychological damage, and so on. We will see later in my discussion of both edge cases this second function proper to belief earning its keep. But it is worth forestalling an objection to this being a proper function of belief at all. It might be thought that putative cases of false beliefs which are nevertheless useful ought not to be characterized as resulting from mechanisms doing what they are supposed to do, but rather as resulting from mechanisms aiming

However, not all accounts of belief which de-prioritize truth and evidence will have this result (see, e.g. Nolfi 2015).

⁷ If it is taken as part of the explanandum of transparency that it holds as a matter of metaphysical necessity, then my explanation will fall short. I have argued against this theoretical constraint elsewhere (2018, §5.3). It might also be thought that cases of non-transparent deliberative beliefs are normatively deficient, and my account cannot explain why this is so. I have argued elsewhere that we can explain why folk may have such intuitions regarding certain epistemic states of affairs. In particular, I have argued that beliefs which concerning epistemic properties and norms are false but adaptive (Sullivan-Bissett 2017a), and I have argued for a biological error-theoretic account of epistemic duties-talk (Sullivan-Bissett 2020).

at truth, and missing. In such cases these false beliefs arise from a function failing to be performed and are merely accidentally useful.

However, there are good reasons to think that this isn't the right way to approach such cases. An oft-cited claim comes from Shelley E. Taylor and Jonathan D. Brown (1994) who have it that unrealistic beliefs about oneself lead to engagement in adaptive behaviours and 'optimal mental health is associated with *unrealistically positive* self-appraisals and beliefs' (see McKay and Dennett 2009: 505–8 for discussion). And, as Kate Nolfi has recently pointed out in her discussion of the optimism bias:

research supplies compelling evidence that we are, as a general rule, more successful in achieving our various ends when our beliefs about ourselves and about our relationship to the world around us are systematically distorted in particular ways. (2018: 192)

Other examples are suggestive of a function proper to belief which is not directed at truth. Some might see our moral beliefs in this way, as adaptive insofar as they facilitate cooperation and social cohesion (Ruse 1986, Joyce 2001), and we might also think of beliefs about epistemic normativity as ones which are adaptive in their making us better at responding to epistemic considerations (Street 2009; Sullivan-Bissett 2017a, 2020). And, as we'll see later, some folk have given adaptationist accounts of self-deceptive attitudes (Trivers 2000, Livingstone-Smith 2014) (§6.3). All of these cases are ones where the cognitive mechanisms responsible are not responding to truth, but are, according to some folk at least, nonetheless producing adaptive beliefs. It would be surprising if all of these cases were ones where things had gone *wrong*, especially given how widespread some of them are, and how they look to contribute positively to inclusive fitness. The existence of these beliefs might look, then, to be in virtue of design not mistake.

Of course, given that deliberative belief formation is characterized by transparency, our mechanisms of belief formation must be operating below the level of consciousness when they aim at something other than truth. Even when mechanisms perform proper function two, we do not distinguish first-personally between beliefs produced in line with proper function one, and those produced in line with proper function two. There is a sense in which we take all our beliefs to be adapted devices, we take them to, in some sense, correspond with, or reflect the environment – we take them all to be true.

5.4 Interim conclusion

Let us take stock. I claimed that the category of *belief* is unified by its essential nature, captured by its motivational role. I identified three features of belief indicative of a link to truth, and argued that the essential nature of belief, however understood, was unable to explain these features. I turned to give a picture of the contingent way beliefs are realized in the actual world, in particular, one which understands them as adapted devices produced by mechanisms with the relational proper function of producing true beliefs. I showed how this function attribution could explain belief's relationship to truth. Finally, I introduced a second function proper to our mechanisms of belief-

production – the production of useful beliefs – in recognition of the fact that belief-production is not exclusively geared towards truth.

Now that my account of belief is on the table, and we have seen how it does in explaining some key features of ordinary belief formation, let us turn to edge cases.

6. Edge cases

How does my view get on with those beliefs which are a bit tricky, which aren't about the colour of lawns or felines on soft furnishings? Here I'll briefly consider two such cases, demonstrating that my view has a great many resources which can be appealed to in order to capture them *as beliefs* (via the modal component), but also explain why they are on the edge at all (via the functional component). My cases will be religious attitudes (§6.2) and self-deceptive attitudes (§6.3)

The commonality among these cases is that the attitudes produced by the mechanisms responsible are often taken to be false,⁸ and they do not look belief-like in various ways. I suggest that non-doxasticism about edge cases is largely motivated by a characterization of a belief which ties it closely to truth, evidence, rationality, etc. Beliefs must be aimed at truth. Beliefs must be evidence-responsive. Beliefs must form a coherent set. These claims are all but taken as truisms. When cases fail to measure up, non-doxasticism threatens (see McCormick 2022: 105–6 for a similar complaint).

I propose rejecting these claims, at least in their modal reading. In releasing the essential nature of belief from the shackles of this strict conception, we can be more nuanced. Many beliefs, maybe most beliefs, have the features these strong characterisations demand. Some don't. But that doesn't disqualify them from doxastic status. In fact, doxastic status comes relatively cheaply. We can appeal to the contingent circumstances of belief's manifestation in the actual world to see why (1) most beliefs meet the strict conception, but also why (2) edge cases do not. We get the explanatory wins of capturing the well-behaved beliefs and the conditions under which they are formed, without giving up on doxasticism about the less well-behaved cases. Belief is cheap. Belief attribution should be tolerant. With the theoretical resources of my approach we can explain the huge range of beliefs proper without giving up on doxasticism or positing in-between states (à la Schwitzgebel 2002, Egan 2008).

The threat of non-doxasticism for edge cases remains, even when we're in austerity mode, insofar as their motivational role appears discordant with that characteristic of belief. These attitudes then, on the face of it at least, might not look to be ones which pass the test set by belief's essential nature. They also raise a more general question about how they fit into the functional account of belief's contingent nature sketched earlier.

There is, of course, a quick way out of the problem where it arises: simply bite the bullet and deny that the target attitude is a belief. This might not even strike one as a theoretical cost or an unpleasant bullet. If one does feel the sting of going non-doxastic, simply shrug one's shoulders and take these attitudes to be casualties of an otherwise plausible and unifying account. Whichever route one takes to non-

⁸ I leave open the possibility of all of these attitudes having true contents. Nothing in the standard characterisations of these attitudes rules that out, although in practice, particularly for self-deception, it might be argued that their mode of formation almost always does.

doxasticism, we would then no longer need to concern ourselves with how these cases fit into my functional account of belief, they are simply outside of that account's remit. However, doxasticism about both of these attitudes is the current orthodoxy, and I'm not in the business of being needlessly revisionary. I will briefly address the alleged belief-discordant motivational role of these attitudes, before assuming for the sake of argument that they are beliefs and so within the remit of my account.

6.1 Malfunction and misfunction

The edge cases I discuss have motivated some authors to advocate a move away from unification in favour of a pluralistic approach to the nature of belief, or indeed, indeterminacy. My way of accommodating such cases is in the spirit of the pluralist recommendation, if not the letter (I retain modal unification). My account can explain why the below cases are ones of belief (via belief's motivational role) and a range of contingencies relating to their biological manifestation can explain some key features which put them on the edge.

To make good on this, we need to say a little more about the distinction between a trait *malfunctioning* and a trait *misfunctioning*. For example, subjects with Holt-Oram syndrome have congenital heart defects caused by mutations in the TBX5 gene with resulting impaired cardiac function. This might be considered a case of *malfunction* (Sullivan-Bissett 2017b: 2512–13). In contrast, a heart which failed to pump blood during its time in a cold box ahead of transplantation would be failing to perform its function, due to being in abnormal conditions (Davies 2000: 33), but it wouldn't be malfunctioning. This would be a case of *misfunction*.

When presented with a candidate belief then, we can say one of three things about it: (1) things have gone well, it is a belief produced by mechanisms performing their proper function, or, things have not gone well and it is either (2) a malfunctioning belief, or (3) a misfunctioning belief. With these options in mind, I turn to my edge cases.

6.2 Religious attitudes

By *religious attitudes*, I mean to refer to those with contents such as *God exists*, *God created the Heavens and the Earth in six days*, *Jesus is the Son of God*, and so on.

Van Leeuwen has argued for non-doxasticism about religious attitudes, understanding them instead as a unique kind of *credence*. He distinguishes factual belief (understood as 'the most mundane, ordinary sense of belief' (2014: 701)) from secondary cognitive attitudes (including imagining, hypothesising, accepting) by appeal to three features necessary to the former, and absent in the latter. Given the account developed in this chapter, I focus only his first feature (practical setting independence), and put aside the other two (cognitive governance and evidential vulnerability).

Van Leeuwen takes it that Donald Davidson is close to correct in characterizing (factual) beliefs as 'mental states that (other things equal) cause actions that satisfy desires, if the beliefs are true' (Davidson 1963, cited in Van Leeuwen 2014: 701). However, he takes it that other attitudes meet this condition, and so it calls for some adjustment to do the work of delineating belief. The first feature of factual belief then is that it is *practical setting independent*, which is to say that it 'guides behaviour in all

practical settings in which [its] content is relevant to the agent's behaviours' (2014: 702). This distinguishes it from other attitudes like *imagining* which is practical setting dependent, guiding behaviors only in some situations. I said earlier (§3) that by enumerating the contexts in which belief plays the motivational role specified, we can distinguish it from other attitudes. For the purposes of this discussion, we can treat my characterization of belief's motivational role and Van Leeuwen's claim that factual beliefs are practical setting independent as picking out roughly the same idea.

Do religious attitudes play the motivational role of belief?⁹ Van Leeuwen says not, and draws on Rita Astuti and Paul L. Harris's (2008) study on members of the Vezo tribe in Madagascar. When asked about the physical and psychological properties of deceased ancestors, those given a religious ritual narrative were more likely to attribute to the deceased psychological properties than those who were given a naturalistic narrative (Van Leeuwen 2014: 706). Van Leeuwen suggests that this can be explained if we take religious attitudes to be *practical setting dependent* and 'becoming deactivated outside the religious-ritual setting' (2014: 706, cf. Ichino 2023: 95–8).

In response, Maarten Boudry and Jerry Coyne argue that religious attitudes may well be 'indeterminate, poorly developed, or partly inconsistent', but this is consistent with them being beliefs (2016: 605). They draw a comparison with a range of other beliefs: *scientific* (folk might endorse the truth of quantum mechanics without understanding what that involves), or *political/economic* (people might report that *there's a recession* without understanding exactly what that means). In such cases we do not deny that these contents are believed. When it comes to attitudes concerning the afterlife in particular, they note that Astuti and Harris's study asked things about the deceased ancestors which fall outside of the concerns of the Vezo tribe. Tribe members care about how the ancestors behave, how best to appease them, how to retrieve their assistance, and so on. Such interests are ones which can tolerate ambiguity when it comes to the psychological profiles of the ancestors. Overall Boudry and Coyne note that:

Given that religious doctrines are often far from coherent, and people resolve inconsistencies by invoking 'mysteries,' we should not be surprised that actual displays of religious belief vary depending on circumstances and priming. (2016: 605)

Boudry and Coyne also offer a range of examples of religious attitudes motivating behaviour in the manner of belief: evangelical Christians building creationist museums, publishing book and articles, and campaigning for the claim that *the earth is 6000 years old* to be taught in schools, Christian scientists choosing to pray for their sick children over seeking medical attention, a jihadist terrorist killing himself and

⁹ I have dropped the 'factual', and so I might be accused of talking past Van Leeuwen. Indeed, he takes *factual belief* to be a term of art, and leans on this when responding to critics to his position (see his (2017) response to Neil Levy (2017)). Might Van Leeuwen have it that religious attitudes are beliefs, but just not *factual* ones? Perhaps. But for our purposes this does not matter. Elsewhere Van Leeuwen has characterised belief simpliciter as the *practical ground* for other attitudes (2009: 219), something which as I have said, can be understood as similar enough to the motivational role essential to belief I picked out earlier. The key point for this discussion is that he denies that this condition is met by religious attitudes.

others in pursuit of paradise, those who go to great financial cost setting up experiments to get a peek at the afterlife, and so on. We are promised that the authors could cite more examples besides (2016: 603–4).¹⁰ For the sake of argument, let us go forward on the basis that religious attitudes do indeed play the motivational role of belief. How do these attitudes fit in my overall framework?

Religious experience, scripture, testimony, philosophical argument, and so on, may well represent an evidence base for religious beliefs. If that's right, then we can understand them as being produced by mechanisms seeking to perform proper function one. Perhaps they are properly tracking the evidence, perhaps they are being misled. However, the fact that religious attitudes are objects of study in philosophy, psychology, and cognitive science, among other disciplines, suggests that we don't have an otherwise unremarkable kind of attitude here, distinguished only by content. *Religious beliefs* are not simply as interesting as say, *geography beliefs* (my belief that *God exists* is of greater theoretical interest than my belief that *Paris is in France*¹¹).

Wherein does the theoretical interest lie? Religious attitudes are often reported as beliefs by their bearers, and yet display features thought to rule them out from that characterization. Most often remarked upon is their relationship to evidence, with many authors arguing that religious attitudes are not supported by evidence, and are impervious to counterevidence. In the previous paragraph I offered some possible sources of evidence for religious beliefs. However, as Anna Ichino has pointed out, *appealing to evidence* should be distinguished from *sensitivity to evidence*. Religious attitudes, according to Ichino:

are related to evidence in the sense that religious people appeal to some evidence to justify them; but the evidence which is appealed to seems to be available only to corroborate, rather than to revise and falsify. Appealing to evidence is not the same as being sensitive to it. (Ichino 2023: 84)

For irresponsiveness to counterevidence, a striking example is found in the behaviour of doomsday cults which do not cease functioning when prophesied dates for the end of the world come and go without incident. Van Leeuwen describes this particular lack of evidential responsiveness as 'astonishing' (2017a: 55). Less extreme cases apparently abound, for example, young earth creationists seem impervious to geological data (2017a: 55).

A subject's religious attitudes may also, *prima facie* at least, conflict with some (other) beliefs the subject holds. As Ichino points out, even if certain religious and

¹⁰ Van Leeuwen responds by pointing out that he was concerned with the attitudes of 'common religious people' in particular, and Boudry and Coyne run together attitudes of fanatics (e.g. suicide bombers) with attitudes of millions of religious folk (2016: 617). The nature of bombers' religious attitudes may well not be informative when we're thinking about common religious people, and so may be put aside. But not all of Boudry and Coyne's examples can be put aside in this way – those publishing books and articles, mounting campaigns to school boards, setting up experiments, and so on, are not *fanatics*, and look much closer to the kind of folk Van Leeuwen's theory is purported to cover.

¹¹ Or even my belief that *Lund is in Norway*. It is not because religious beliefs are false (if they are) that they are of interest. Of course their contents, insofar as they're metaphysically bold, might make them especially interesting, but I take it that is not the only reason they are an independent topic of study.

scientific beliefs were reconcilable, what matters is that religious folk often hold attitudes that *seem* inconsistent, that they themselves *take to be* inconsistent, without attempting to overcome the inconsistency (Ichino 2023: 88–9).

Van Leeuwen and Ichino opt for non-doxasticism in light of the above considerations (among others), a position recommended to them by their respective conceptions of belief. On my view of course, belief and evidence are not so tightly aligned, and where they are, this is secured by their contingent biological manifestation, in particular, by mechanisms of belief production seeking to produce true beliefs. In light of their special theoretical interest, possibly arising from their relationship to evidence, we might think that this is not how religious beliefs should be accommodated by my account. We might instead take these attitudes to be produced by mechanisms of belief production seeking to perform proper function two, the production of useful beliefs.¹² This is in keeping with recent work in cognitive science on the genesis and role of religious beliefs and explanations. For example, Telli Davoodi and Lombrozo ran three studies to investigate whether scientific and religious explanations play different roles (i.e. epistemic versus non epistemic). Their third study is of particular interest here. Participants were asked to provide explanations of one of three existential questions¹³ with either a focus on epistemic characteristics (evidence, logic) or non-epistemic characteristics (emotional comfort, peace of mind) (2022: 1210). Scientific explanations were more prevalent than religious explanations when participants were given the epistemic instruction, and the reverse was true when participants were given the non-epistemic instruction (2022: 1213). As they note in their general discussion:

Beliefs can serve different purposes, [...] If religious explanations are perceived (by some) as best suited to serving nonepistemic needs, they are likely to arise and persist in response to such needs. (2022: 1215)

The suggestion is that the cognitive mechanisms responsible for the production of religious beliefs may not be operating in the service of accurate representation or other epistemic goals. If this were right, this would fit the broad architecture I have outlined.

Van Leeuwen argues that religious attitudes ‘have different etiologies and different cognitive and behavioural effects’ from factual belief (2014: 698). My ontically austere approach to belief has the resources to accept this without giving up doxasticism or positing a new attitude of *religious credence*. We can grant a class of *factual beliefs* with particular properties, and a particular etiology (produced by mechanisms seeking to perform proper function one), but also grant religious attitudes doxastic status (secured by their meeting belief’s motivational role), with their particular features and *raison d’être* accommodated by their genesis being traced to mechanisms seeking to perform proper function two.

¹² In the next section I will suggest the possibility of self-deceptive beliefs being produced by such mechanisms. Georges Rey (2007) has argued that religious attitudes are the result of self-deception, which, if right, can be subsumed under the explanations in the next section.

¹³ *How did the universe come to exist? Why is there suffering in the world? What happens after we die?* (Davoodi and Lombrozo 2022: 1201, Table 1).

6.3 Self-deception

Let us begin with a characterization of self-deception which is a theory-neutral as possible, drawing on what Van Leeuwen (2007) recognizes as three ‘strands of consensus’. Take the self-deceptive attitude to be *that p*. First, a motivational component is involved in the production of the attitude (in the interests of space we’ll take this to be a *desire that p*¹⁴). Second, there is information available to the person deceived which would justify believing what they ought to believe, for simplicity, not-*p*. Third, self-deception results in a truth-evaluable cognitive attitude.

Is the attitude arising from self-deception a belief? The question is harder to answer for this particular edge case because self-deception is more elusive, insofar as the literature relies on imagined cases,¹⁵ where authors are free to describe things in a manner friendly to their views. This is different from religious attitudes which can often enough be delineated by their contents (cf. Van Leeuwen 2014: 701), as well as other potential edge cases like delusion for which we have (albeit highly imperfect) diagnostic criteria, and implicit biases for which we have a range of implicit measures (see Sullivan-Bissett 2023, §2 for an overview). We can point to cases of religious and delusional attitudes *and observe them*. We can then disagree about whether *x* piece of behaviour or *y* piece of reasoning, etc. was belief-concordant. For self-deception, we rely on intuition pumping vignettes. For example, Robert Audi in his argument for non-doxasticism and in favour of his avowal view uses the example of Ann, a woman dying from cancer. In asking whether she believes the content to which she is self-deceived he says we can attribute to her the undesired belief that *she will die* (because she enquires into funeral arrangements, has a reasonable medical knowledge, and adjusts her will) (1982: 139). Does she also believe the content about which she is deceived (*that she will recover*)? Audi says no: ‘She says this, but “actions speak louder than words,” and her overall behavior will not support the attribution of this belief’ (1982: 139).

Van Leeuwen on the other hand, in responding to Audi, uses the cases of a businesswoman who takes out loans for a failing business (she is self-deceived that *the business will become successful again*), and the case of the Air Florida Flight 90 where the pilot was warned that the instrument panel was displaying problematic readings but nevertheless allowed take off (he was self-deceived that *the plane is safe enough to fly*). After criticizing the avowal view Van Leeuwen argues that his cases are ones which meet the conditions on belief (the attitudes govern other cognitive attitudes in the manner of belief, and are the default for action (2007: 435)).

If there is nothing in the three strands of consensus that speaks to whether the product of self-deception is belief, our target is a moving one. It is the additional details of usually imagined cases that are going to prop up a verdict regarding belief

¹⁴ This is a case of *wishful* self-deception, where the desired state of affairs is the content of the self-deceptive attitude. Other forms of self-deception are *wilful*, where the desire is to *hold the belief that p* (rather than for the state of affairs *p* to obtain) and *dreadful*, where the desire is for not-*p* and the agent’s self-deceptive attitude is that *p* (Mele calls this *twisted* self-deception, 2001: 4) For discussion see Van Leeuwen (2007).

¹⁵ Even in real life examples like the Air Florida Flight 90 to be discussed below, it’s not clear we have the tools to properly determine whether we have a case of self-deception. We have to infer the relevant architecture paradigmatic of the phenomenon (see Trivers and Newton (1982) who analyse this case in such terms).

status. To be clear, that philosophers rely on imagined vignettes in their discussions is not objectionable. Rather, the problem is that in this particular discussion, the conditions which need to be met for an imagined case to qualify as self-deception are sufficiently thin so as to require us to embellish those cases with details relevant to our question. That is, nothing in Van Leeuwen's strands of consensus speaks to whether the product of self-deception guides one's action in the manner of belief. So this feature is an optional add-on when one is describing an imaginary case.

And so, on my view, self-deception generates *belief* if its products play the motivational role specified earlier. Do they? This is not straightforwardly ascertained. We can equally recognize that Audi's Ann's attitude does not play the role, and that Van Leeuwen's businesswoman's attitude does. Both cases are ones of self-deception insofar as they meet the three strands of consensus. Perhaps we need some further strands which might rule out one set of cases in favour of another, or perhaps we should accept that self-deception sometimes results in belief, and other times not. It is, of course, beyond the scope of this piece to resolve this methodological quirk in the literature. Instead, let us for the sake of argument, and to cover our bases, take the products of self-deception to be the business of an account of belief, and see how my account does if that is right.

I turn now to two broad approaches to the architecture of self-deception (intentionalism and non-intentionalism), and then various biological accounts of self-deception (adaptationist and non-adaptationist). I will show that all of these approaches are consistent with the cognitive architecture proposed in my overall account.

Intentionalist accounts have it that self-deception is a kind of action whose goal is to get the self-deceiver to form a belief that she takes to be false. It is in intentionalism that we find the most clear analogy with interpersonal deception. Of course, many people take it to be impossible to intend to get oneself to believe something one believes is false, and our doxastic deliberation being transparent to considerations of truth straightforwardly prevents doing this in a conscious manner, even if one could form the appropriate intention. In light of concerns of this kind, intentionalists have appealed to *partitioning*, either temporal or psychological. On a temporal story, I simply arrange the world such that I will be tricked in the future, ensuring that I do not remember the arrangements. If this is the right account, then the products of this process can be understood as arising from the performance of our mechanisms for belief production seeking to perform proper function one. And missing. The oddness of self-deceptive beliefs would come not from the psychological mechanisms responsible for the resulting belief, but from the purposeful arranging of the environment to mislead those mechanisms.

Another kind of partitioning is psychological, where self-deceivers hold contradictory beliefs. The cognitive system responsible for the deception is hidden from the agent, so she is free to believe that *p*, whilst also, the belief that not-*p* is hidden from the conscious self. Or, for those intentionalists who do not think that the deceptive part of the architecture harbours the belief that not-*p*, at the very least it will harbour the unconscious intentions to form the self-deceptive belief. If this is the correct picture, the resulting attitudes are best understood as produced by mechanisms seeking to perform proper function two. Indeed, we might think that a

context in which there is significant motivation for a particular state of affairs to obtain is abnormal for the performance of proper function one. Instead, mechanisms of belief production seek to produce a belief which would be protective.

What about non-intentionalist approaches? Very roughly, a non-intentionalist has it that self-deception is a result of motivational biases. For example, Alfred Mele has given detailed characterizations of indirect ways that a desire that p may have an influence over the process of belief formation resulting in the belief that p (2009: 56–9). Broadly, a desire that p may influence selective attention to particular evidence and the adjusting of confidence thresholds for belief based on evidence. Again, if this is right, then self-deceptive beliefs can be understood as those produced by mechanisms aiming at beliefs that are useful.¹⁶ Let us turn now to biological approaches.

There are broadly two ways of thinking about a biological explanation of self-deception, one which casts it as an *adaptation*, and the other which casts it as a *spandrel*.¹⁷ Robert Trivers argues that the capacity for self-deception makes an organism a better interpersonal deceiver. An appeal to the mechanisms of ordinary deception is made, something which readily lends itself to an adaptationist account. Deception between individuals may generate self-deception, as a mechanism for preventing detection by others of ongoing deception (Trivers 2000: 115). I can deceive my adversary better if I am unaware that I am being deceptive, because this leaves avenues for the detection of deception closed. On this approach then, the conscious mind is construed as a social front, ‘maintained to deceive others who more readily attend to its manifestations than those of the actor’s unconscious mind’ (Trivers 2000: 115).

David Livingstone-Smith has also developed an adaptationist model, which he pitches as an alternative to intentionalist and non-intentionalist approaches. He criticizes Mele on the grounds that his approach doesn’t allow us to distinguish self-deception from other kinds of motivated reasoning (for example, wishful thinking). He notes that intentionalists about self-deception can do this work (the key difference between self-deception and other motivated reasoning is that the former has intentional character) (2014: 185). However, intentionalists have their own problems, for example, having to appeal to a deceiving homunculus. On his preferred teleofunctional non-intentionalist theory, self-deception has a biological purpose. His view captures what is good about intentionalism (it is able to distinguish self-deception from other motivated reasoning), but avoids the problems by appealing to a sub-personal mechanism whose job it is to selectively prevent the organism’s

¹⁶ As Eric Funkhouser has pointed out, ‘[m]any hot biases and cases of wishful thinking are not purposive – at least, they do not have the purpose of distorting belief’ (2019: 188). It might be thought then that a non-intentionalist account like Mele’s cannot be translated into a functional analysis in which mechanisms of belief formation are seeking to perform proper function two. But the two are in fact consistent. I have understood the second function as the production of *useful beliefs* (although not useful as an approximation to truth). Those mechanisms are operative in contexts abnormal for the fulfilment of proper function one. But they need not be understood as seeking to *distort* belief, but rather understood as seeking to form beliefs with a particular feature, that of being useful.

¹⁷ Stephen J. Gould and Richard C. Lewontin use the term *spandrel* to pick out phenotypic traits that are structural by-products rather than adaptations. They use the example of a dome mounted on rounded arches, which results in tapering triangular spaces underneath the dome which are ‘necessary architectural byproducts of mounting a dome on rounded arches’ (1979: 581).

representational apparatus from performing its proper function of accurate representation (2014: 191–2).

If either of these approaches were right, they can both be subsumed under my framework. In particular, we can understand self-deceptive beliefs as generated by our mechanisms of belief-production seeking to perform proper function two. Here those mechanisms would not only be seeking to produce a useful belief, but would be doing so *in opposition* to the function of true belief production. That is, we could understand the mechanisms in this context as not simply doing their job of useful belief production, but rather *disrupting the process of true belief production* in the service of doing their job (cf. fn. 17).

Let us turn to a spandrel account of self-deception. According to Van Leeuwen, the capacity for self-deception has arisen as a spandrel upon three other (adaptive) features of mind. The first is the ‘sting’ which accompanies desires; felt when anticipation of their non-fulfilment is present, which can help motivate the achieving of one’s goals. The second is the ability to selectively attend to evidence, which is essential to any creature with finite cognitive resources and interests. The third is the inclination toward harm avoidance, the function of which is to keep us away from situations in which we may be harmed (2008: 199). Van Leeuwen claims that these three features can work together to give rise to self-deceptive belief, for example: I desire that p and so consequently, upon anticipation that $\sim p$, I feel the sting (feature one). Such anticipation may be supported by evidence that $\sim p$. Discomfort is felt when I attend to the evidence that $\sim p$ but fades when I do not – particularly when attending to the evidence that p . Consequently I attend to the evidence that p and self-deception ensues (features two and three). I focus attention on the evidence that p and so come to believe that p despite the total evidence pointing to $\sim p$ (2008: 198–9).

We can frame this account in my overall picture in roughly the same way as we understood adaptionist approaches. Mechanisms of belief production operating in the service of proper function two have in their gift the ability to facilitate selective attention to evidence in certain circumstances (i.e. when particular desires and anticipated harms are present). We can understand these circumstances as ones abnormal for the performance of proper function one. In facilitating selective attention to evidence, a useful belief is generated, and, in some cases, that will be properly characterized as a *self-deceptive belief*. This honours the idea of the capacity of self-deception as a spandrel, and simply absorbs the features of mind on which it is parasitic into the operation of proper function two.

In sum then, whatever story we prefer regarding the nature of self-deception (intentional versus non-intentional) and the biology of self-deception (adaption or spandrel) we can see how it finds a coherent home in my overall framework. For intentionalism with temporal partitioning, we have a case of malfunctioning belief, where mechanisms of belief production are seeking to perform proper function one, but have been indirectly manipulated by the agent’s arrangement of the environment. In all other cases we have beliefs produced by mechanisms performing their proper function (the production of useful beliefs). This applies to intentionalism with psychological partitioning, non-intentionalism, and across biological approaches.

6.4 Bringing cases back from the edge

I have briefly discussed two cases which may be taken to constitute a challenge to an account of belief. I opened the chapter by noting that in seeking a unifying property, one had better make it cheap. Imposing strict conditions on belief leaves us with a more sparsely populated category than we should want, excluding the kinds of cases we have considered in this section. My view is able to bring these cases back from the edge. They are straightforwardly tokens of belief like any other, insofar as they meet the motivational condition. We can explain why these cases have been cast to the edge by appeal to the contingent biological circumstances of their formation.

7. In defence of ontic austerity

Let us bring together the various threads of this piece which constitute my defence of ontic austerity. There are broadly three virtues to be gained from taking belief to be modally cheap, with additional characteristics as explainable by appeal to its contingent manifestation in biological creatures.

The first virtue is *unification*. My view allows us to capture beliefs in a single category by appeal to a single modal condition (§3). This allows us to capture garden variety beliefs as well as edge cases (§6).

The second virtue is *explanatory power*: my picture allows us to explain belief's relationship to truth as secured by the contingent biological circumstances of some beliefs (§5.2), as well being consistent with, and insightful to, various things that might be said about edge cases (§6.)

Finally, my view demonstrates the virtue of parsimony. It is better if we do not have to revise our folk psychological understanding of cognition, or its components.¹⁸ To accommodate edge cases, some philosophers have proposed new categories (Egan's bimagination (2008) to accommodate delusion, Tamar Gendler's alief (2008) to accommodate self-deception, Van Leeuwen's religious credence (2014) to accommodate religious attitudes). These authors are forced to such moves in part because they work with an intolerant conception of belief, one which is ill-equipped to accommodate the quirks of some of its tokens. My view does better. All of these cases can be properly recognized as beliefs, and those features which have led some to non-doxasticism are better thought of as arising from their particular biological manifestation.

8. Conclusions

The account developed here opens avenues for a better understanding of belief's various guises—from garden-variety beliefs about shopping lists, to beliefs embedding strange metaphysical commitments, right through to beliefs under the spell of motivational influences.. We have seen that it is attention to the modally local question concerning belief in the actual world which promises to better illuminate various phenomena associated with belief and its formation. Bulking up the modal character of belief to include various epistemic notions only closes off explanatory gains in our picture of the operation of doxastic deliberation, and leads us having to

¹⁸ This principle has wide support. For example, Levy has it that 'we ought to avoid multiplying mental states unnecessarily, we need to ensure that we are postulating exotic states and processes only when they are truly needed' (Levy 2016: 9).

accept non-doxasticism for various states of mind which guide action and are sincerely avowed as beliefs. Let us enjoy the cake of unification, whilst eating the cake of explanation. With a bonus sprinkling of parsimony. Let us be ontically austere.

Acknowledgements

I acknowledge the support of the Arts and Humanities Research Council (*Deluded by Experience*, grant no. AH/T013486/1). Thank you to Federico Bongiorno, Nikk Effingham, Jonathan Jong, Michael Rush, and an anonymous reviewer for comments on an earlier version of this paper. Special thanks to Miriam McCormick for discussing the material with me.

References

- Armstrong, David M. 1973: *Belief, Truth and Knowledge*. London: Cambridge University Press.
- Astuti, Rita and Harris, Paul L. 2008: 'Understanding Mortality and the Life of the Ancestors in Rural Madagascar'. *Cognitive Science*. Vol. 32, no. 4, pp. 713–40.
- Audi, Robert 1982: 'Self-deception, Action and Will'. *Erkenntnis*. Vol. 18, no. 2, pp. 133–58.
- Boudry, Maarten and Coyne, Jerry 2016: 'Disbelief in Belief: On The Cognitive Status of Supernatural Beliefs'. *Philosophical Psychology*. Vol. 29, no. 4, pp. 601–15.
- Bykvist, Krister and Hattiangadi, Anandi 2013: 'Belief, Truth, and Blindspots'. In Chan, Timothy (ed.) *The Aim of Belief*. Oxford: Oxford University Press, pp. 100–22.
- Davies, Paul 2000: 'Malfunctions'. *Biology and Philosophy*. Vol. 15, pp. 19–38.
- Davidson, Donald 1963: 'Actions, Reasons, and Causes'. *Journal of Philosophy*. Vol. 60, pp. 685–700.
- Davoodi, Telli and Lombrozo, Tania 2022: 'Explaining the Existential: Scientific and Religious Explanations Play Different Functional Roles'. *Journal of Experimental Psychology*. Vol. 151, no. 5, pp. 1199–1218.
- Dokic, Jerome and Engel, Pascal 2002: *Frank Ramsey: Truth and Success*. London: Routledge.
- Dular, Nicole and Fortier, Nikki 2021: 'Epistemic Reasons, Transparency, and Evolutionary Debunking'. *Philosophia*. Vol. 49, pp. 1455–73.
- Egan, Andy 2008: 'Imagination, Delusion, and Self-Deception'. In Bayne, Tim and Fernandez, Jorge (eds.) *Delusion and Self-deception: Affective and Motivational Influences on Belief Formation*. New York: Psychology Press, pp. 263–80.
- Funkhouser, Eric 2019: *Self-deception*. Routledge.

Forthcoming in Jong, Jonathan, and Schwitzgebel, Eric (eds.) *The Nature of Belief*. Oxford University Press. [Please cite final version.]

Gendler, Tamar Szabó 2008: 'Alief and Belief'. *Journal of Philosophy*. Vol. 105, no. 10, pp. 634–63.

Glüer, Kathrin and Wikforss, Asa 2013: 'Aiming at Truth: On the Role of Belief'. *Teorema*. Vol. 42, no. 3, pp. 137–62.

Gould, Stephen C. and Lewontin, Richard C. 1979: 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme'. *Proceedings of the Royal Society of London*. Vol. 205, pp. 581–98.

Ichino, Anna 2019: 'Imagination and Belief in Action'. *Philosophia*. Vol. 47, no. 5, pp. 1517–34.

Ichino, Anna 2023: 'Religious Imaginings'. In Sullivan-Bissett, Ema (ed.) *Belief, Imagination, and Delusion*. New York: Oxford University Press, pp. 81–106.

Joyce, Richard 2001: *The Myth of Morality*. Cambridge: Cambridge University Press.

Kind, Amy 2023: 'Contrast or Continuum? Belief and Imagination'. In Sullivan-Bissett, Ema (ed.) *Belief, Imagination, and Delusion*. New York: Oxford University Press, pp. 42–59.

Levy, Neil 2016: 'Have I Turned the Stove off? Explaining Everyday Anxiety'. *Philosophers' Imprint*. Vol. 16, no. 3, pp. 1–10.

Levy, Neil 2017: 'Religious Beliefs are Factual Beliefs: Content Does Not Correlate With Context Sensitivity'. *Cognition*. Vol. 161, pp. 109–16.

Livingstone-Smith, David 2014: 'Self-Deception: A Teleofunctional Approach'. *Philosophia*. Vol. 42: pp. 181–99.

Mandelbaum, Eric 2016: 'Attitude, Inference, Association: On the Propositional Structure of Implicit Bias'. *Noûs*. Vol. 50, no. 3, pp. 629–58.

Mayo, Bernard 1964: 'Belief and Constraint'. *Proceedings of the Aristotelian Society*. Vol. 64, pp. 139–156.

McHugh, Conor 2012: 'Beliefs and Aims'. *Philosophical Studies*. Vol. 160, no. 3, pp. 425–39.

McHugh, Conor 2013: 'The Illusion of Exclusivity'. *European Journal of Philosophy*. Vol. 23, no. 4, pp. 1117–36.

McCormick, Miriam 2015: *Believing Against the evidence: Agency and the Ethics of Belief*. Abingdon, Oxon: Routledge.

McCormick, Miriam 2022: 'Belief as Emotion'. *Philosophical Issues*. Vol. 32, no. 1, pp. 104–119.

McKay, Ryan and Dennett, Daniel 2009: 'The Evolution of Misbelief'. *Behavioral and Brain Sciences*. Vol. 32, pp. 493–561.

Mele, Alfred 2001: *Self-deception Unmasked*. Princeton: Princeton University Press.

Mele, Alfred 2009: 'Self-Deception and Delusions'. In Bayne, Tim and Fernandez, Jordi (eds.) *Delusion and Self-Deception*. New York and London, Taylor and Francis, pp. 55–69.

Forthcoming in Jong, Jonathan, and Schwitzgebel, Eric (eds.) *The Nature of Belief*. Oxford University Press. [Please cite final version.]

Millikan, Ruth 1984: *Language, Thought, and Other Biological Categories*.

Millikan, Ruth 1989: 'In Defense of Proper Functions'. *Philosophy of Science*. Vol. 56, no. 2, pp. 288–302.

Millikan, Ruth 1993: 'Propensities, Exaptations, and the Brain'. In Millikan, Ruth (ed.), *White Queen Psychology and Other Essays for Alice*. Cambridge: MIT Press, pp. 31–50.

Nolfi, Kate 2015: 'How to be a Normativist about the Nature of Belief'. *Pacific Philosophical Quarterly*. Vol. 96, pp. 181–204.

Nolfi, Kate 2018: 'Why Only Evidential Considerations Can Justify Belief'. In McHugh, Conor, Way, Jonathan, and Whiting, Daniel *Normativity: Epistemic and Practical*. Oxford: Oxford University Press, pp. 179–99.

Noordhof, Paul 2001: 'Believe What You Want'. *Proceedings of the Aristotelian Society*, Vol. 101, no. 1, pp. 247–65.

O'Brien, Lucy 2005: 'Imagination and the Motivational View of Belief'. *Analysis*. Vol. 65, no. 1, pp. 55–62.

Papineau, David 1987: *Reality and Representation*. Oxford: Basil Blackwell Limited.

Rey, Georges 2007: 'Meta-atheism: Religious Avowal as Self-Deception'. In Antony, Louise (ed.) *Philosophers Without Gods: Meditations on Atheism and the Secular Life*. Oxford: Oxford University Press. pp. 243–65.

Ruse, Michael 1986: 'Evolutionary Ethics: A Phoenix Arisen'. *Zygon*. Vol. 21, no. 1, pp. 95–112.

Schellenberg, Susanna 2013: 'Belief and Desire in Imagination and Immersion'. *Journal of Philosophy*. Vol. 110, no. 9, pp. 497–517.

Schwitzgebel, Eric 2002: 'A Phenomenal, Dispositional Account of Belief'. *Noûs*. Vol. 36, no. 2, pp. 259–75.

Shah, Nishi 2003: 'How Truth Governs Belief'. *The Philosophical Review*. Vol. 112, no. 4, pp. 447–82.

Shah, Nishi and Velleman, J. David 2005: 'Doxastic Deliberation'. *The Philosophical Review*, Vol. 11, no. 4, pp. 497–534.

Sinhababu, Neil 2013: 'Distinguishing Belief and Imagination'. *Pacific Philosophical Quarterly*. Vol. 94, pp. 152–65.

Stalnaker, Robert C. 1984: *Inquiry*. Cambridge University Press.

Steglich-Petersen, Asbjørn 2006: 'No Norm Needed: On the Aim of Belief'. *The Philosophical Quarterly*, Vol. 56, no. 225, pp. 499–516.

Street, Sharon 2009: 'Evolution and the Normativity of Epistemic Reasons'. *Canadian Journal of Philosophy*. Vol. 39, pp. 213–48.

Forthcoming in Jong, Jonathan, and Schwitzgebel, Eric (eds.) *The Nature of Belief*. Oxford University Press. [Please cite final version.]

Sullivan-Bissett, Ema 2017a: 'Biological Function and Epistemic Normativity'. *Philosophical Explorations*. Vol. 20, no. 1, pp. 94–110.

Sullivan-Bissett, Ema 2017b: 'Aims and Exclusivity'. *European Journal of Philosophy*. Vol. 25, no. 3, pp. 721–31.

Sullivan-Bissett, Ema 2017c: 'Transparency Defended'. *Syndicate Philosophy*, symposium on Miriam McCormick's *Believing Against the Evidence: Agency and the Ethics of Belief*.

Sullivan-Bissett, Ema 2018: 'Explaining Doxastic Transparency: Aim, Norm, or Function?' *Synthese*. Vol. 195, no. 8, pp. 3453–76.

Sullivan-Bissett, Ema 2020: 'We Are Like American Robins'. In Stapleford, Scott and McCain, Kevin (eds.) *Epistemic Duties: New Arguments, New Angles*. New York: Routledge, pp. 94–110.

Sullivan-Bissett, Ema 2023: 'Implicit Bias and Processing'. In Thompson, Robert (ed.) *The Routledge Handbook of Philosophy and Implicit Cognition*. Routledge, pp. 115–26.

Sullivan-Bissett, Ema and Noordhof, Paul 2020: 'The Transparent Failure of Norms to Keep Up Standards of Belief'. *Philosophical Studies*. Vol. 177, pp. 1213–27.

Taylor, Shelley E. and Brown, Jonathan D. 1994: 'Positive Illusions and Well-being Revisited: Separating Fact from Fiction'. *Psychological Bulletin*. Vol. 116, no. 1, pp. 21–7.

Trivers, Robert 2000: 'The Elements of a Scientific Theory of Self-deception'. *Annals of the New York Academy of Sciences*. Vol. 907, pp. 114–31.

Trivers, Robert and Newton, Huey P. 1982: 'The Crash of Flight 90: Doomed by Self-deception?' *Science Digest* (November), 66–67, 111.

Van Leeuwen, Neil 2007: 'The Product of Self-Deception'. *Erkenntnis*. Vol. 67, pp. 419–37.

Van Leeuwen, Neil 2008: 'Finite Rational Self-Deceivers'. *Philosophical Studies*. Vol. 139, pp. 191–208.

Van Leeuwen, Neil 2009: 'The Motivational Role of Belief'. *Philosophical Papers*. Vol. 38, no. 2, pp. 219–46.

Van Leeuwen, Neil 2014: 'Religious Credence is not Factual Belief'. *Cognition*. Vol. 133, no. 3, pp. 698–715.

Van Leeuwen, Neil 2016: 'Beyond Fakers and Fanatics: A Reply'. *Philosophical Psychology*. Vol. 29, no. 4, pp. 616–21.

Van Leeuwen, Neil 2017a: 'Do Religious Beliefs Respond to Evidence?' *Philosophical Explorations*. Vol. 20, pp. 52–72.

Van Leeuwen, Neil 2017b: 'Two Paradigms for Religious Representation: The Physicist and the Playground (a Reply to Levy)'. *Cognition*. Vol. 164, pp. 206–11.

Forthcoming in Jong, Jonathan, and Schwitzgebel, Eric (eds.) *The Nature of Belief*. Oxford University Press. [Please cite final version.]

Van Leeuwen, Neil and Lombrozo, Tania 2023: 'The Puzzle of Belief'. *Cognitive Science*. Vol. 47, e13245, pp. 1-5.

Velleman, J. David 2000: *The Possibility of Practical Reason*. Oxford: Oxford University Press.

Williams, Bernard 1973: 'Deciding to Believe'. In his *Problems of the Self*. Cambridge: Cambridge University Press, pp. 136-51.