

Biological Function and Epistemic Normativity

Sullivan-Bissett, Ema

DOI:

[10.1080/13869795.2017.1287296](https://doi.org/10.1080/13869795.2017.1287296)

License:

Creative Commons: Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Sullivan-Bissett, E 2017, 'Biological Function and Epistemic Normativity', *Philosophical Explorations*, vol. 20, no. Sup 1, pp. 94-110. <https://doi.org/10.1080/13869795.2017.1287296>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.


Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.



Philosophical Explorations

An International Journal for the Philosophy of Mind and Action

ISSN: 1386-9795 (Print) 1741-5918 (Online) Journal homepage: <http://www.tandfonline.com/loi/rpex20>

Biological function and epistemic normativity

Ema Sullivan-Bissett

To cite this article: Ema Sullivan-Bissett (2017) Biological function and epistemic normativity, *Philosophical Explorations*, 20:sup1, 94-110, DOI: [10.1080/13869795.2017.1287296](https://doi.org/10.1080/13869795.2017.1287296)

To link to this article: <http://dx.doi.org/10.1080/13869795.2017.1287296>



© 2017 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 11 Apr 2017.



Submit your article to this journal [↗](#)



Article views: 385



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 1 View citing articles [↗](#)

Full Terms & Conditions of access and use can be found at
<http://www.tandfonline.com/action/journalInformation?journalCode=rpex20>

Biological function and epistemic normativity

Ema Sullivan-Bissett*

Philosophy Department, University of Birmingham, Birmingham, UK

(Received 22 November 2016; final version received 5 December 2016)

I give a biological account of epistemic normativity. My account explains the sense in which it is true that belief is subject to a standard of correctness, and reduces epistemic norms to there being doxastic strategies which guide how best to meet that standard. Additionally, I give an explanation of the mistakes we make in our epistemic discourse, understood as either taking epistemic properties and norms to be *sui generis* and irreducible, and/or as failing to recognise the reductive base of epistemic normativity. This explanation will appeal to the claim that the beliefs which constitute our epistemic discourse are false but adaptive, and are the outcome of a non-truth-tracking process. The opponents of my position are philosophers who take epistemic normativity not to be reducible in this way, and to involve *sui generis* properties and norms governing belief. The aim of the paper is to show that epistemic normativity can be explained by appeal to the biological functions of our mechanisms of belief-production.

Keywords: epistemic normativity; epistemic norms; biological function; belief; truth; correctness

1. Preliminaries

Some of those interested in the nature of morality have been moved by considerations of strangeness, a conversation made most famous by Mackie (1977). Since then, some philosophers (e.g. Mackie 1977; Ruse 1986, 1993; Joyce 2001) have given error accounts of our moral discourse and the norms and properties such discourse cites. Mackie directed his arguments towards moral properties and discourse, but thought that his conclusions generalised to, for example, the aesthetic case. If that is right we should also take them to generalise to the epistemic case, since our epistemic discourse equally commits us to metaphysically strange properties, “entities that are hard to square with a naturalistic world-view” (Olson 2011, 88).

Mackie argued that “there do not exist entities or relations of a certain kind, objective values or requirements” (1977, 17). This point has two prongs. First, objective values if they existed would be strange entities, qualities, or relations. Consider the property of *being morally forbidden*. A little reflection on this reveals that “too much is being asked of the world – there is simply nothing that is forbidden in the specifically moral sense of the word” (Joyce 2001, iv). Second, if we were aware of these strange entities, qualities, or relations, the way in which we were aware of them would be via a correspondingly

*Email: e.l.sullivan-bissett@bham.ac.uk

strange faculty, different from the way in which we are aware of everything else (Mackie 1977, 38).

Our epistemic discourse also commits us to objective requirement, insofar as we take epistemic reasons to be *categorical* reasons. Note that I am not basing an argument on Mackie's and others' positions here. Rather, I am pointing out that there are error-theoretic accounts given of other discourses, and so perhaps we can give an error-theoretic account of our epistemic discourse too, especially since our epistemic discourse has the features Mackie identifies as problematic in the moral case, namely, the postulation of strange properties and the notion of objective requirement.

For the purposes of this paper I assume a motivational account of belief according to which it is necessary to believe that it plays some motivational role. I adopt the role offered by Lucy O'Brien, namely that belief is an attitude that "by itself, and relative to a fixed background of desires, disposes the subject to behave in ways that would promote the satisfaction of his desires if its contents were true" (O'Brien 2005, 56). I take features of belief linking it to truth to be contingent ones. What makes something a *belief* then, is not its being truth directed (or any nearby claim), but whether or not it plays some motivational role. If the reader does not find the condition I have put on belief appealing, she should feel free to replace it with her favourite motivational role of belief and input that as the background to my account. I have argued elsewhere that building belief's relationship to truth into the very nature of belief (as advocated by teleological and normative accounts, for example), is problematic insofar as it cannot do the explanatory work advocates claim for it (Sullivan-Bissett 2017). Additionally, it has recently been argued that nor can such accounts (so-called, *epistemic constitutivist* accounts) ground epistemic normativity (Côté-Bouchard 2016). So I think it is worth seeing how far we can get with an alternative account of belief's nature in an explanation of epistemic normativity. For those readers dissatisfied with the starting assumption of a motivational account of belief, I invite them to place the paper on a conditional: if the motivational account of belief is right, then epistemic normativity could be explained as my view proposes.

In ascribing biological functions to our mechanisms of belief-production, I will adopt a historical account of biological according to which, very roughly, the function of some trait is to do whatever ancestors of that trait did which got them selected.¹ For it to be appropriate to ascribe a biological function to our mechanisms of belief-production then, they need to have been selected for. There are several reasons to think that this is the case, though space constraints prevent outlining them here (interested readers should see Millikan 1993c, 49, and discussion in Fales 1996). Of note though is that, given my view that beliefs are necessarily motivational devices, the very nature of belief means that adaptive pressures will be in play. The formation of beliefs produces changes to the way a subject navigates within and responds to its environment, and so the mechanisms which produce them will be sensitive to adaptive pressures.

2. Truth is not the only end

It has been argued in many places that our mechanisms of belief-production have as their biological proper function the production of true beliefs (for representative views, see Papi-neau 1987; Millikan 1993a, 1993b, 1993c). Without dwelling on the reasons in favour of this claim, note that beliefs which are true are more likely to dispose a creature to act in ways which will satisfy its desires, be they for food, warmth, sex, and so on. The claim that true beliefs are adaptive has been taken to be obvious by many philosophers, Quine for example claimed that "creatures inveterately wrong in their inductions have a pathetic

but praiseworthy tendency to die before reproducing their kind” (Quine 1985, 39). Of course, true beliefs are not always adaptive (see, for example, Stich 1990; Stephens 2001), but such exceptions should not deter us from accepting the plausible claim that they usually are (Cowie 2014, 4007; Street 2009, 235).

For the sake of exposition, let us understand this claim as one which says that the *only* function of our belief-producing mechanisms is to produce true beliefs. In this section I present some cases of usually false belief in which the mechanisms which produce them are functioning as they should. I do this to warm the reader up to the idea of there being a biological function proper to the mechanisms which produce our beliefs, which is other than to track truth.

The existence of false beliefs does not present a problem for the claim that our mechanisms for belief-production have as their only biological function the production of true beliefs. This is because claiming a function for a set of cognitive mechanisms says nothing about how often (or even if) this function will be performed. As Ruth Millikan puts the point:

a description of the biological functions of the cognitive systems will in no way resemble a catalogue of psychological laws. It is certainly no psychological law, for example, that our beliefs are true, though it is a (teleo)function of our belief-fixing systems to fix true beliefs. (Millikan 1993a, 177)

Functions can fail to be performed, and something possesses a function because *in certain key moments*, the performance of it contributed to the reproductive success of its bearers. (Compare with the case of sperm: it is no objection to the claim that sperm have the proper function of fertilising ova that they rarely perform this function – “[m]ost never find an ovum and have to call it quits” [Millikan 1984, 34].)

What constitutes a problem for the *only* claim are those cases of false belief which seem to have been produced by mechanisms doing what they were *supposed to do*, that is, producing adaptive beliefs, which are usually false. If there are cases in which our mechanisms for belief-production are functioning properly (that is, doing what ancestral tokens of that type were selected for doing) when they produce such beliefs, this falsifies the claim that the only function of these mechanisms is to produce true beliefs.

Candidate cases are those of beliefs produced via self-enhancement bias, partiality bias (different doxastic treatment of one’s friends over strangers), and self-deception, to name just a few.² Cases like these show that the mechanisms responsible for belief-production are not solely geared, in all cases, towards truth. Admittedly this is a little quick, but note that such cases represent a large set of beliefs held by ordinary people, which exhibit biases which are not truth-directed, and so suggest (if not clearly demonstrate) that we need to recognise another function proper to our belief-producing mechanisms, one which can accommodate these beliefs. I suggest that it is the function of producing beliefs which have the role of being useful in self-organisation. By *organising beliefs*, I will mean those beliefs which facilitate self-organisation, maintain self-esteem, avoid psychological damage, and so on. These organising beliefs do not have their utility in virtue of approximating to truth, but rather in their assisting the effective functioning of the believer.

Adopting Millikan’s notion of “Normal” (note the capitalisation), used to designate a normative historical (not statistical) sense of normality: the various mechanisms which give rise to organising beliefs do so in circumstances abNormal for the production of true beliefs. I will label the production of true beliefs *proper function one*, and the production of organising beliefs in the sense just specified *proper function two*.

There are two worries with my picture here, which I will answer before moving on. The first is that I have taken cases of, for example, self-deceptive beliefs, to be produced by mechanisms functioning to produce *organising* beliefs operating as they should, and not by mechanisms functioning to produce *true* beliefs and missing their target. An alternative account of the production of these beliefs might be that they are the outcome of a malfunction of the general capacity to produce true beliefs. Why prefer my model of their production? Here I appeal to Ryan McKay and Daniel Dennett's (2009) distinction between two types of *misbelief*: those resulting from a malfunction of the mechanisms of belief-production, and those resulting from the Normal functioning of those mechanisms. Misbeliefs of the first type are the result of mechanisms failing to perform the function of producing true beliefs. Misbeliefs of the second type are the result of mechanisms of belief-production functioning as they should. The worry then about my description of the production of beliefs resulting from various doxastic biases is that they are actually instances of the first kind of misbelief, and not the second.

To diffuse this worry I note that there are good reasons to think that these beliefs are not malfunctioning ones, and are rather the result of general evolved capacities: they have a certain psychological role to play. For example, recent empirical work has shown that "optimal mental health is associated with *unrealistically positive* self-appraisals and beliefs" (see McKay and Dennett 2009, 505–508 for discussion). Such unrealistic beliefs about oneself lead to engagement in adaptive behaviours (Taylor and Brown 1994). Although McKay and Dennett's focus is on positive illusions as Normally produced beliefs (and not cases of malfunction), I think the other candidates I identified for organising beliefs introduced above (and indeed, below) can be given a similar story. The existence of these beliefs then is not in virtue of mistake, but of design.

The second worry is related to the way I have carved up the proper functions of our mechanisms for belief-production; perhaps neither of my proposed function ascriptions are correct. It is not that there are two functions: the production of true beliefs, and the production of organising beliefs, there is rather just one function: the production of *useful* beliefs. However, to capitulate to this worry would only be to re-describe the machinery I am suggesting is in place. We could move to *usefulness* to capture what our mechanisms for belief-production are up to, but only insofar as we can move to *usefulness* to capture what all traits with biological functions are up to. We ascribe to the chameleon's mechanisms of skin pigmentation the function of producing a camouflaging skin pattern, this is useful insofar as it prevents predation. At a more coarse-grained level of description, the function of these mechanisms is to produce *useful* skin patterns. This point is a trivial one. My description of two proper functions as belonging to our mechanisms of belief-production is supposed to map onto two sets of Normal conditions involved in the production of beliefs: Normal conditions for the production of beliefs which are true, and Normal conditions for the production of beliefs which are organising. Of course both of these functions produce useful items, but usefulness will be part of the story for all functional items. That is just bound up in the nature of function ascription.

A more sophisticated version of this worry might be the following: if we grant that belief is the kind of thing which plays some motivational role parsed in terms of satisfying desires when its contents are true, then why not identify the biological function of our mechanisms of belief-production as producing items which are suited to play this role in such a way as to satisfy desires? This would capture those contexts in which they do so *by producing items with true contents*, and those contexts in which they do so *by producing items with organising contents*.³ In response to this, I note that I have no objection to putting things in these terms. If the biological function of our mechanisms of belief-production

were to produce devices well suited to lead to desire satisfaction, we would still need to specify two sets of Normal conditions for the performance of that function: those conditions in which beliefs lead to desire satisfaction by being true, and those conditions in which they do so by being organising. To repeat: talk in terms of different functions is intended to map onto these different sets of Normal conditions. Moving to the claim that the function of our mechanisms of belief-production is to produce devices which lead to desire satisfaction will still require distinguishing these contexts.

I have given some examples of beliefs which are usually false, but which are nevertheless produced by mechanisms of belief-production functioning as they should. This will be important later in my explanation of the intuitive pull of explanations which cite *sui generis* epistemic properties and norms or which at the very least fail to note the reductive base of epistemic normativity.

3. Epistemic normativity

I will give an explanation of epistemic normativity by appeal both to proper function one and proper function two. Insofar as I appeal to proper function two, an outcome of my account is that at least some of our beliefs about epistemic normativity are false but organising. I will use the term *epistemic normativity* to pick out the conjunction of two claims. First, belief has a standard of correctness according to which true beliefs are correct and false beliefs are incorrect.⁴ Second, there are *sui generis* epistemic norms which govern belief.

Looking first to 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). Though my *imaginings* or *supposings* will have contents which are true or false, they are nevertheless not appropriately judged as correct or incorrect. Correctness conditions then follow not only from the propositional content of a state, but also from the state itself.

The standard of correctness for belief is that a belief is correct if and only if it is true. "Correct" is not synonymous with "true," but is thought to be a normative notion which attaches not to the proposition believed, but to the attitude or act of believing. Following Krister Bykvist and Anandi Hattiangadi, I will understand belief's standard of correctness as non-normative, understanding correctness in this way is in line with common usage:

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. (Bykvist and Hattiangadi 2013, 103)

The point here is a simple one: some standards generate an ought, some standards do not, and this depends on which other conditions obtain. For example, judging that "driving as a woman in Saudi Arabia contravenes conventional standards is not to think a woman in Saudi Arabia ought not to drive" (Bykvist and Hattiangadi 2013, 103). Here we have a standard which many people would not endorse, and so it is perhaps easier to see the difference between a standard being in place and there being something normative about that standard. In other cases, the difference may be obscured by our already endorsing the standard.

I appeal to Bykvist and Hattiangadi here in a pre-emptive spirit. Those who think standards of correctness necessarily entail there being something normative about that standard, will already take my account as developed below to be wrongheaded, since it is in this

weaker sense of standard that I will give an explanation of the standard of correctness for belief. It is important then to be clear that there is a difference between a standard being in place and there being something normative about that standard which generates an *ought*. I will later explain why people are prone to think that there is a *sui generis* normativity attached to the standard of belief, or at least, why people fail to recognise the biological basis of this standard. I will also suggest that the standard of correctness for belief is not always truth (see Sections 4.3 and 4.4).

Turning to epistemic norms which are claimed to govern belief formation. The most obvious and discussed of such norms is grounded in belief's standard of correctness: Pascal Engel claims that it "is generally agreed that the normative dimension in belief is its dimension of correctness, and that the norm for belief, if there is such a norm, is that a belief is correct if and only if it is true" (Engel 2007, 180).

Other kinds of norms include those of evidence (a belief is correct if it rests upon sufficient evidence), knowledge (a belief is correct if and only if it aims at knowledge), and rationality (a belief is correct if and only if it is rational) (Engel 2007, 181). These norms are taken to govern only belief – it is generally inappropriate to say of my imaginings or supposings that they are rational, irrational, justified, unjustified, and so on (and where it is not inappropriate, such judgements are grounded on the downstream effects of the attitude). These epistemic norms have been understood as categorical ones, that is, as ones where our obligations to comply with them are those "to which the practical benefits of beliefs are not relevant. They are obligations that arise from a purely impartial and disinterested perspective" (Feldman 1988, 236).

What I take to require an explanation then are the following two claims.

(EN1) Beliefs have truth as their standard of correctness.

(EN2) There are *sui generis* categorical epistemic norms.

The explanatory burden is to give an account of why (EN1) and (EN2) hold or at least, why we think that they do (when they do not). My account will pave a middle way. I will argue that there is a sense in which (EN1) and (EN2) are true, but not in the way that they are typically taken to be. My explanation will go via the claim that our mechanisms of belief-production have the function of producing true beliefs, and this is the only sense in which true beliefs are correct and false beliefs are incorrect (EN1). Given belief's standard of correctness, there are strategies for doxastic regulation which facilitate the formation of beliefs in line with this standard. So there are not *sui generis* norms governing belief as (EN2) states, but rather these norms are reducible to doxastic strategies which facilitate the meeting of belief's standard of correctness.

4. Reducible and irreducible epistemic normativity

There are at least two ways to understand claims (EN1) and (EN2), relating to how we understand the normativity involved. The opponent to my view is she who takes claims (EN1) and (EN2) at face value, as involving *sui generis* irreducible epistemic normativity, or weaker, as not being reducible to biological function. There are two kinds of explanation important to my discussion: one which cites *sui generis* epistemic normativity or is silent with respect to the reductive base of epistemic normativity, and my preferred explanation which is a biological reductionist explanation. I am going to refer to the positions which give these two explanations as *irreducible epistemic normativity*, and *reducible epistemic normativity*.

The way in which I think that (EN1) and (EN2) are true then, is the way which I have called *reducible epistemic normativity*, and this can be explained by appeal to proper function one. However, many of us find it intuitive⁵ that (EN1) and (EN2) are to be understood more in line with what I have called *irreducible epistemic normativity*. So as well as giving an account of why reducible epistemic normativity is true, I will also give an explanation of why we make the mistake of thinking that irreducible epistemic normativity is true. This mistake can be explained by appeal to proper function two.⁶

4.1. *Explaining reducible epistemic normativity*

Before giving my explanation of reducible epistemic normativity, a digression is needed to note my use of “categorical” in this paper, which is important for being clear about the commitments of my view. I said earlier in my explication of epistemic normativity that epistemic norms have been taken to be categorical, and I cashed out what this means in terms of independence from agents’ goals or desires. However, a little more is needed, since categoricity has been understood along at least two dimensions. Often these two dimensions are treated as equivalent, or at least, as not requiring separation, as we will see in explications from Kant and Chase B. Wrenn below. To see the various parts of my view, these dimensions should be kept separate.

The first dimension is independence from the goals or desires of agents which the rule or norm binds. This is contrasted with hypothetical norms or rules as ones which are binding only insofar as they are means to some other end. Kant explicates categoricity in this way, noting that “[t]he categorical imperative would be that which represented an action as necessary of itself without reference to another end, that is, as objectively necessary” (Kant 1785/2005, [414]/74). A categorical imperative “would express a reason for acting which was unconditional in the sense of not being contingent upon any present desire of the agent to whose satisfaction the recommended action would contribute as a means” (Mackie 1977, 29).

Another dimension of categoricity though, is its binding all rational agents. Again, Kant speaks of categorical imperatives governing all subjects who are “practically determinable by reason” (Kant 1785/2005, [414]/74). And Wrenn claims that a categorical rule is one which “binds *everyone*, regardless of their particular ends” (Wrenn 2004, 279, my emphasis).

Often these two dimensions are kept together (as in the previous quote from Wrenn) – a rule or norm which is categorical in the sense of applying to agents independently of their aims or desires, is also one which is categorical in the sense of applying to *all* agents of a given type, say, all *rational* agents. My view is one which understands so-called epistemic norms as doxastic strategies which facilitate the meeting of belief’s standard of correctness. These strategies do not represent categorical imperatives in either sense. Though they might be good strategies regardless of the interests of agents (“believe in line with perceptual experience” is a strategy which will facilitate believing in line with belief’s standard of correctness, even if I have no desire to have true beliefs), they do not generate any claims about how the agent *ought* to believe, and so following them is not “objectively necessary.” These strategies are also not categorical in the sense of applying to all rational agents, or all believers, and that is because not all beliefs will have truth as their standard of correctness (see Sections 4.3 and 4.4).

I now move to explaining (EN1) and (EN2). My explanation of (EN1) is the following: our belief-producing mechanisms have as one of their proper functions the production of true beliefs, and this provides the only sense in which true beliefs are correct and false

beliefs are incorrect. This is in the same way that there is a standard from which a chameleon's skin pattern can deviate determined by the environment which makes it a correct or incorrect skin pattern. In the case of belief, their contents can vary with respect to their meeting the correctness conditions laid down by biology. When a belief is true, the mechanisms which produced it have performed their function, when a belief is false the mechanisms have failed to perform their function (or stronger, have malfunctioned). It is in this sense only that beliefs have a standard of correctness.

Beliefs produced by mechanisms functioning to produce true beliefs have the derived proper function of being true. What does that say about what the individual *believer* must do? *Absolutely nothing*. As Papineau notes, it "is a vulgar, and indeed dangerous, error to infer, from the premise that X has been biologically designed to do y, that in some sense X *ought* to y" (Papineau 1999, 21, fn. 5). Our mechanisms of belief-production being designed to produce true beliefs does not mean that they *ought* to, nor does it mean that we, at the agent-level, *ought* to facilitate the meeting of this standard. This is important, because our epistemic practice suggests that we do feel obligations in this area, and yet, if the standard of correctness for belief is grounded in something which does not generate obligations, what are we to say of our epistemic practice? The idea I will develop is the following one: our epistemic practice is not tracking any real obligation on us to believe truly, rather, the derived biological function of many of our beliefs is to be true, and our epistemic practice facilitates the meeting of that *biologically grounded* standard.

My explanation of (EN2) is more a denial of the phenomenon so characterised. Epistemic norms are reducible to doxastic strategies which facilitate the meeting of belief's standard of correctness. These strategies are mischaracterised as categorical epistemic norms which bind our doxastic practice. The kinds of strategies I have in mind are those like "treat normal perceptual experience as *prima facie* veridical, honor logical inferences, and employ the inductive methods in empirical enquiry" (Leiter 2015, 75). These are strategies which "facilitate successful navigation of the world and prediction of the future course of experience" (Leiter 2015, 75). Why do we take these strategies to be more heavy-duty, and to involve normative force? At least part of the answer is that creatures who tend towards judging that considerations of truth represent *normative reasons* for belief will be more likely to arrive at true beliefs about their environment, which, as I suggested earlier, is adaptive (Street 2009, 234). Similar things can be said about regulating one's belief formation by considerations regarding what would be a *justified* belief, or a *rational* belief, and so on.

One way of thinking about my view is to note the ways in which I depart from my opponents. I disagree with the proponent of irreducible epistemic normativity on three grounds: I disagree that there are *irreducible epistemic* norms, I disagree with the global extension of those norms (I do not think all *believers* are governed by such norms), and I disagree with local extension of norms governing belief (I do not think all of *our beliefs* are governed by such norms). I will take each point in turn.

4.2. Irreducibility

On my view there are no irreducible epistemic properties or norms. My explanation of (EN1) is one which appeals to proper function one of our mechanisms for belief-production. This proper function provides the sense in which truth is the standard of correctness for beliefs. My explanation of (EN2) goes via an appeal to this biological standard, and notes that there are doxastic strategies which facilitate meeting it. Epistemic norms then are reducible to strategies one can adopt which make it more likely that one will have true beliefs, beliefs produced by mechanisms which are successful in performing proper

function one. My opponent takes belief's standard of correctness to be *epistemic* in kind, and takes there to be categorical epistemic norms which bind our doxastic practice.

4.3. *Global extension*

My view is that (EN1) and (EN2) do not hold for all believers. I take the standard of correctness for belief and the doxastic strategies which facilitate meeting it, to be a result of the biological history of our mechanisms of belief-production, and so this standard and the strategies which facilitate its being met do not hold for all believers. I do not think, for example, that it is already determined that Swampman's beliefs are correct when true, and if that is right, strategies like "believe in line with perceptual experience" are not going to be ones which facilitate meeting the standard of correctness for Swampman's beliefs (if they have one). This is because Swampman does not have a biological history.⁷ For a rational creature with a different (biological) history, my opponent will take belief's standard of correctness and epistemic norms to apply to them. For my account, whether this creature's beliefs are correct when true, and so whether certain doxastic strategies will facilitate meeting that standard, is an empirical question, in particular, one which may turn on that creature's biology or other relevant historical factors.

4.4. *Local extension*

Third, and most controversially, I disagree with my opponent about the local extension of (EN1) and (EN2). I am seeking to explain (EN1) and (EN2) for the large subset of our beliefs which are produced by mechanisms performing (or trying to perform) proper function one: the production of true beliefs. But, as I argued earlier, it would be remiss of us not to note that many of our beliefs are produced by mechanisms seeking to produce *organising* beliefs, whose usefulness is not in their approximating to truth, but in facilitating the self-organisation of the believer. Those beliefs which are produced by mechanisms of belief-production functioning to produce *organising* beliefs, do not come under the standard of truth for belief. This means that some false beliefs are *correct* (insofar as they are the result of mechanisms functioning properly in line with proper function two), and some false beliefs are *incorrect* (insofar as they are the result of mechanisms functioning improperly in line with proper function one). This is a consequence of locating the correctness of belief in the proper functions of the mechanisms which produce them. This means that, not only do (EN1) and (EN2) not apply to all believers, they also do not apply to all of the beliefs of humans.

5. Explaining (away) irreducible epistemic normativity

In the previous section I gave my explanation of the sense in which (EN1) and (EN2) are true, and noted how my view departs from proponents of irreducible epistemic normativity. However, this is an explanation which, I suspect, will strike many as unintuitive. If you already agree that epistemic normativity (as identified as constituted by (EN1) and (EN2)) can be given a biological explanation, and you have no intuitions in contrast with that, and already take the epistemic discourse you engage in to involve systematic error, I have completed my journey with you. I suspect though that most readers will not be persuaded by explanatory power alone. As both Mackie and Richard Joyce note in their discussion of the moral case, in claiming that some discourse commits us to systematic error, one ought to be able to give an account of why we commit these mistakes (Mackie

1977, 17–18; Joyce 2001, 135). As such, it falls upon me now to give consideration to how we think about the claims in (EN1) and (EN2), identify the mistakes we make in our epistemic discourse, and explain why we do so.⁸

5.1. Identifying the error(s)

The account one offers of the mistakes we make in our epistemic discourse will vary, depending on exactly what the mistakes are. I think there are two options. The first is to understand our epistemic discourse as one which posits *sui generis* irreducible epistemic properties and norms, the existence of which my account denies. The second is to understand our epistemic discourse as one which is incomplete in its failing to recognise the reductive base of epistemic normativity. This would be to understand epistemic discourse merely as one which does not present as tracking some biological standard, but nor does it commit us to *sui generis* irreducible epistemic normativity. I will offer an account which can explain why we make either or both of these mistakes.

Does our epistemic discourse involve the positing of *sui generis* irreducible epistemic properties and norms, and beliefs in line with this? When we say things like “it is correct to believe that that *Neil Armstrong walked on the moon* because it is true that *Neil Armstrong walked on the moon*,” do we take ourselves to be appealing to a special kind of normativity? Put in this way, perhaps not. However, those philosophers engaged in conceptual analysis of belief, who take belief to be subject to an *epistemic* standard which is not merely a biological one, might support this characterisation of our epistemic discourse. If their analysis of the concept of belief is right, then this will map onto common practice. Consider also those who take the practice of inquiry very seriously, people who are so driven to find truths – such people might be correctly characterised as feeling strong *epistemic* obligations, in line with the characterisation of epistemic discourse involving appeal to *sui generis* irreducible norms and properties.⁹

Perhaps our epistemic discourse does not involve anything so sophisticated, but rather, just fails to reflect the truth about the reductive base of epistemic normativity. So when we say things like “it is correct to believe that *Neil Armstrong walked on the moon* because it is true that *Neil Armstrong walked on the moon*,” we are not citing *sui generis* epistemic properties or norms, we are just failing to recognise the only sense in which it is true that believing this would be correct – that doing so would be to believe in line with belief’s biologically grounded standard of correctness.

Note that if this is the right characterisation of the error inherent in epistemic discourse, it is not just that we are ignorant of the reductive base. We also make the mistake of thinking that belief’s standard of correctness (however it is grounded), generates claims about how we *ought* to believe. And, as I noted earlier, that is an error. Our mechanisms of belief-production having the biological function of producing true beliefs entails nothing about how we *ought* to form beliefs. So even if our epistemic discourse does not commit us to *sui generis* epistemic properties, at the very least it commits us to ought-claims about what to believe, or to thinking that we have *reasons* to believe truly, regardless of our interests. On my view, those claims are false.

Ought-claims alone are not kryptonite to the naturalist, since naturalists can perfectly well accommodate *hypothetical* oughts. But our epistemic discourse does not commit us to these harmless imperatives. We think that we can have epistemic reasons to believe propositions even if “one’s believing those propositions holds no promise of advancing any goal which one actually possesses” (Kelly 2003, 630). We think that we can have epistemic reasons to believe propositions which are “not contingent upon whether [we] care about

believing what is true” (Cuneo 2007, 59). Even if our mistake with respect to the nature of epistemic normativity is in our not recognising its reductive (biological) base, we also make judgements about having *reasons* to believe which are not hypothetical in kind.¹⁰

We need not resolve exactly what we commit ourselves to in our epistemic discourse, since what I will say below, with appropriate amendments, can explain either of these characterisations. In laying them out though, I have sought to capture whatever it is such discourse involves, and what they have in common, namely, that they take the standard of correctness for belief to be normative, and they take there to be categorical epistemic norms. So though it is up for grabs where the error lies, at the very least, when we make judgements about epistemic reasons, “we have a strong sense that we are tracking facts that are robustly independent of us” (Street 2009, 243). I agree with Sharon Street that we are right to think this, but I depart from the irreducible epistemic normativity theorist, and indeed from Street, with respect to the way in which it is right. What matters for the moment though is that whatever we are tracking is not normative, and is rather reducible to biological function.

Some readers at this point may be on board, and may see the explanatory burden our epistemic discourse places on a biological reductionist account. Others might accuse me of over-intellectualising our discourse, and might say that there is no such burden. If these others are right, I have no argument with them.¹¹ However, here I answer those who think that in giving a biological reductionist account of epistemic normativity I leave something unexplained.

5.2. Explaining the error(s)

We have seen two ways in which our epistemic discourse might be in error: by positing *sui generis* epistemic normativity, and/or by failing to acknowledge the biological nature of epistemic normativity. Both errors mistakenly take belief’s standard of correctness and the strategies which might facilitate our meeting it, to be prescriptive and categorical.

My strategy is to give a biological error-theoretic explanation of why we make these mistakes, by appeal to the claim that doing so is adaptive. I am interested in explaining why we have those beliefs which are part of our epistemic discourse, even though they are false.¹² For the purpose of exposition, call these *irreducible epistemic beliefs* (a little less cumbersome than “beliefs about irreducible epistemic normativity”). In short, my claim is that having true beliefs is adaptive, and believing in irreducible epistemic normativity facilitates the production of true beliefs. Given this, we should expect that believers would be equipped with domain-specific-regulating mechanisms specifically tailored to produce these irreducible epistemic beliefs.

Wherever we identify the error in epistemic discourse, many of our beliefs thereof are false. Those beliefs about *sui generis* irreducible epistemic properties and norms express propositions which are uniformly false. Those beliefs which fail to reflect the source of epistemic normativity as biological normativity are, at the very least, incomplete. An explanation can be given for why we have such beliefs by appeal to the biological usefulness of them. Our having irreducible epistemic beliefs helps facilitate the meeting of belief’s biologically grounded standard of correctness. In offering an evolutionary explanation which cites a process as non-truth-tracking, I am not saying that the contents of the target beliefs can be explained in evolutionary terms – that *particular beliefs* were selected. Rather, I am accounting for our tendency to have such beliefs by claiming that what has been selected for are “cognitive mechanisms that entail dispositions to form certain primitive belief-like representations in certain environments” (Clarke-Doane 2012, 318).

It is in the following way that irreducible epistemic beliefs could be biologically useful, and it is because of their usefulness, that the mechanisms which produce them may have proliferated. Recall, such mechanisms function to produce organising beliefs, which are useful not as an approximation to truth, but rather useful with respect to facilitating the effective self-organisation of the believer. In order to make our mechanisms of belief-production more likely to perform their proper function of producing true beliefs, our biological history has ensured we have thoughts about correct, incorrect, rational, irrational, justified, unjustified (and so on) belief, where we take such evaluations to involve *sui generis* normativity, or at the very least, we do not take them to be reducible to biology. We hold our beliefs up to *epistemic* standards; we feel we ought to have *sufficient* or *good* evidence for our beliefs and that others *ought* to apply the same standards to their beliefs. We have evolved dispositions which incline us to judge our beliefs and the beliefs of others in these ways. We have an interest in the beliefs of other people: we use them as information sources, and the success of their actions will often depend on their beliefs being true or rational or justified (Wrenn 2004, 284–285). Our feeling it appropriate to hold our beliefs to epistemic norms – to standards of production – is a way of making it more likely that our beliefs have true contents.

Having true beliefs, that is, having beliefs which reflect the biological nature of epistemic normativity, would result in fewer true beliefs than we would have were we to have false beliefs about epistemic normativity (as, I say, we do). If we recognised that the only sense in which true belief is correct belief is the biological sense, we would be less inclined to form beliefs in accordance with what we identify as epistemic norms, that is, doxastic strategies which facilitate the meeting of belief's biologically grounded standard of correctness. Just as in Joyce's parallel account of moral beliefs, these "feelings of 'inescapable requirement' will, in certain circumstances, serve reproductive fitness more effectively" (Joyce 2001, 140) than true beliefs about the nature of epistemic normativity.

This account assumes that we can exercise some kind of control over our beliefs. This is a harmless assumption, since it is one we find in our very practice of epistemic evaluation. As Kate Nolfi notes, though epistemic evaluation does function to mark whether a believer meets some epistemic standard, it also is "evaluation that has directive or instructive import" (Nolfi 2014, 99, fn. 1). The appropriateness of epistemic evaluation is explained in part by our being able to exercise some kind of doxastic control – and our epistemic evaluative practices presuppose this (Nolfi 2014, 109, fn. 19).¹³ Now then, I move on to show how our falsely judging that epistemic normativity involves *sui generis* properties or norms, or our simply not representing epistemic normativity as reducible to biology, is adaptive.

Compare the case of belief-forming practices to our reproductive practices. The function of our reproductive organs is to reproduce. We recognise this and yet many of us choose not to reproduce. We do not allow the standards laid down by biology to motivate us inescapably to act in accordance with them. Though we do not take there to be *sui generis* irreducible categorical procreation norms, there is a mechanism in place which helps facilitate our following that biological imperative (i.e. sexual desire; and/or so-called *broodiness*). We might expect to find some mechanism in the epistemic case present to encourage us to follow the standards for belief laid down by biology. I think our epistemic discourse is just such a mechanism.

Biology has laid down truth as the standard of correctness for belief, that is, we Normally believe truly (in circumstances Normal for the performance of proper function one). There are doxastic practices which facilitate our beliefs meeting that standard. Recognising that there are standards laid down by biology with respect to our belief-forming practices, does not in any way secure conforming to those standards. If we had true beliefs about the

nature of epistemic normativity, we may well be less motivated to form beliefs in epistemically ideal ways, so far as this is psychologically possible, and also might be less disposed to judge the epistemic behaviour of others. Our having beliefs which reflect the true nature of epistemic normativity only guarantees that we recognise in a weak sense, that belief has a standard of correctness and there are ways to go about meeting that standard. What is not guaranteed by the presence of these features is that one believes that *p* on *sufficient* or *reasonable* evidence, or that we feel motivated to respond to putative epistemic norms.

Knowing that the standards and norms for some activity are only biological in kind might allow for reflective distance; one might feel less motivated to adhere to some standard, if that standard is merely a biological one.¹⁴ Having beliefs at the wrong strength with respect to the normativity involved in belief might make it psychologically more difficult not to engage in the kinds of behaviours and belief-forming practices that our epistemic discourse commits us to. That is, such discourse encourages us to exercise doxastic control in line with what we take to be *epistemic* standards and norms.

In addition, my explanation of reducible epistemic normativity speaks against the truth of a stronger reading of (EN1) and (EN2) (which has irreducible epistemic beliefs coming out as true). Note that the claims of reducible epistemic normativity and irreducible epistemic normativity are contraries. To see this, let us see how far we can get when we imagine their both being true. Consider the opposing explanations of (EN1): belief's standard of correctness would be reducible to biology, but also be part of the very concept of belief. This would introduce overdetermination, which is presumably best avoided. Now the explanations of EN2: on my view epistemic norms are reducible to doxastic strategies which we identify as *sui generis* norms guiding our epistemic practice. My opponent though takes there to be these *sui generis* norms which are not reducible as I claim, and given these contradictory claims, we cannot both be right. The falsity of irreducible epistemic normativity follows from the truth of my account of reducible epistemic normativity. And then all that remains is to explain why we nevertheless are under the illusion of irreducible epistemic normativity, which has been the focus of this section.

A more ambitious claim, for which I have not argued, is that epistemic normativity can *only* be explained by appeal to biology. One way to support that claim would be to argue that our irreducible epistemic beliefs are not a good guide for theorising about the nature of epistemic normativity (given their proposed etiology), and so do not provide us with evidence against the biological account given here, and then, perhaps, we have no reason to believe that an explanation in irreducible terms is required. This is a more ambitious project than showing that the biological account has the resources to accommodate the explananda, and is something I do not undertake here.

6. A note on expressivism

One way to develop the account that I have been outlining, is to “go expressivist” about what I have understood to be irreducible epistemic *beliefs*.¹⁵ So claims like “Glen’s belief is irrational,” or “Katie has an epistemic reason to believe that *p*, even though she does not care about forming a belief about *p*,” are not ones which give voice to some *belief* the subject has, but rather are merely expressions of some sentiment, commitment, or other non-cognitive mental state (Bar-On and Sias 2013, 699).

There are two reasons one might be attracted to such a position, given what I have said in the paper. First, considerations of metaphysical strangeness abound when we think about the commitments we make in having attitudes about irreducible epistemic normativity.

Rejecting the claim that such attitudes are in the business of reporting some state of the world makes that strangeness disappear (Bar-On and Sias 2013, 701). Second, expressivist proposals in general have been motivated, at least in part, by their ability “to capture an apparent connection between sincere utterances of claims in that discourse and non-neutrality on the part of the speaker in terms of attitude and motivation” (Bar-On and Sias 2013, 700). We have seen that something like this connection is in play with the target phenomenon here – utterances in epistemic discourse serve the purpose of regulating the belief-forming practices of oneself and others. Adopting expressivism about irreducible epistemic attitudes might capture this feature of them.

For those attracted to such a view, I note that adopting it in the context of what has been argued in this paper would only call for some fairly minor adjustments to the overall account, after all, my proposal as it stands and expressivism about the target phenomenon are in agreement on ontological matters. The adjustment would be that instead of giving an account of why we are disposed to have certain beliefs (even though they are false), the evolutionary explanation would take as its target our having certain sentiments, or commitments. A fuller discussion of this way of developing the account is beyond the scope of the current paper, I just note then that such a development is a possible one.

7. Conclusions

I gave a biological account of epistemic normativity, which explained the sense in which it is true that belief is subject to a standard of correctness, and reduced *sui generis* epistemic norms to doxastic strategies prescribed by our epistemic discourse which facilitate the meeting of belief’s biologically grounded standard. I also gave an explanation for why we have beliefs in line with irreducible epistemic normativity, which went via the claim that having such beliefs is adaptive. I conclude that a biological account of belief has the resources to accommodate belief’s standard of correctness and the epistemic practices which govern belief formation and maintenance.

Acknowledgements

Thanks to audiences at the Mind and Reason research group at the University of York, the Reasons and Norms Research Seminar at the Open University, and the False but Useful Beliefs workshop in London, for helpful discussion on earlier versions of this paper. Thanks also to Anneli Jefferson, Paul Noordhof, Kathy Puddifoot, Jon Robson, Kirk Surgener, and two referees for comments and discussion which improved the paper.

Disclosure statement

No potential conflict of interest was reported by the author.

Funding

I acknowledge the support of the European Research Council [grant agreement 616358] for funding the research of which this is a part.

Notes

1. Versions of this so-called *strong etiological* account can be found in Millikan (1984, 1989), Neander (1991a, 1991b), Papineau (1993), and Sullivan-Bissett (2016). Weaker etiological

accounts which do not require *selection* of the functionally characterised trait but only that past tokens of it contributed to organismic fitness can be found in Buller (1998) and Price (1995, 2001). I think my account could house a weaker etiological theory of function, but for the purposes of this paper, and because my view is that the strong etiological account is correct, I adopt a stronger version.

2. Indeed, evolutionary accounts have already been given of at least some of these phenomena. For example, the capacity to self-deceive has been explained as an evolutionary adaptation (Trivers 2000, 2011, 2013) and as an evolutionary spandrel (Van Leeuwen 2007, 2008).
3. I am grateful to an anonymous referee for this more sophisticated version of the worry.
4. Note that belief's standard of correctness is not a demarcating feature of belief, as this standard is shared by guesses (Owens 2003).
5. For those who do not have the intuition that (EN1) and (EN2) ought to be understood in these terms, I do not have the burden of providing an error theory of that intuition.
6. I should note that the way I have mapped the terrain does not carve things along the line between what Christopher Cowie calls *intrinsicists* and *instrumentalists* about epistemic normativity. According to the first position, epistemic normativity is *sui generis* normativity, so for example, we have reason to believe in line with the evidence because of some "*brutely epistemic normative truth*" (Cowie 2014, 4003). Obviously, this is a view which my position opposes. According to the second position, epistemic normativity is derivative of practical normativity, so for example, we have reason to believe in line with the evidence because by doing so we are more likely to fulfill our goals and satisfy our desires (Cowie 2014, 4003). It might be less obvious how my view is different from this one. I take it this is the kind of position David Papineau has in mind when he argues that there is no distinctive epistemic normativity, and that the rightness or wrongness of believing truly or falsely is traceable to moral, personal, or aesthetic normativity (Papineau 2013). My view is different from this, since I claim that the only sense in which a true belief is *correct* is the sense in which the mechanisms which produced it have performed their biological proper function. So though both myself and epistemic instrumentalists are reductionists, the way in which we are so means that the positions should be kept separate.
7. The case of Swampman might not be an ideal one, since there is room for saying that Swampman is not a rational agent, and thus the dimension of categoricity which is to do with application to all rational agents will not apply here. If that is right, then take a rational alien being. In my view, it is an open question whether (EN1) and (EN2) apply to the beliefs of such a being.
8. There is a substantial literature on the coherence of error theories, which I cannot engage with here. If it turns out that error theories cannot be true, the work here will, at the very least, require serious adjustment. However, there are good reasons to think that general purpose arguments against error theories in general have not yet gotten a hold, and that the prospects for error theory may well have been underestimated (see, for example, Daly and Liggins 2010), and so it is legitimate, I think, to offer a theory in these terms, at least for now.
9. Thanks to Keith Allen for discussion on this.
10. Even the instrumentalist about epistemic normativity recognises that she needs to tell a story about the putative (but non-actual) categoricity of epistemic reasons (e.g. Olson 2011, 86; Papineau 1999, 24).
11. Though see Joyce (2001, 141) for an argument that humans engage in talk of categorical reasons, even if they do not use the correct philosophical jargon.
12. The work I am doing here is merely descriptive, I make no claims about what we ought to do as believers. A sensible naturalised epistemology would look something like the one Papineau develops according to which "you should think of yourself as a system for generating true beliefs. You want to be as reliable such a system as possible. So you should consider ways of redesigning the system, and should implement those that promise an improvement" (Papineau 1987, 135).
13. I agree with most philosophers that we cannot exercise any *direct* control over our beliefs. However, my account of the usefulness of irreducible epistemic beliefs does not require the truth of doxastic voluntarism, it requires only that we can exercise doxastic control *in some sense*. Whatever account one prefers of our ability to exercise doxastic control, this will be congenial to what I want to say about the usefulness of our false beliefs about epistemic normativity. For example, the Immediate Causal Impact account has it that a believer exercises doxastic control when she is caused to regulate her beliefs in some particular circumstances by her

judgements regarding how she ought to regulate her beliefs in those particular circumstances (Nolfi 2014, 97). The Disposition Regulation Account of doxastic control has it that a believer's judgements about how she ought to believe do not directly impact her practice, but rather shape her cognitive character, thus having a deferred impact on her belief-forming practices, by bringing "the way in which she believes closer in line with her conception of the relevant epistemic ideal" (Nolfi 2014, 108). Whichever of these accounts is correct, we can see how judgements of epistemic evaluation can impact the regulation of our beliefs. This occurs either directly via those judgements (Immediate Causal Impact account), or indirectly via those judgements shaping our cognitive characters (Disposition Regulation Account).

14. Of course, if I persuade the reader that epistemic normativity can be given a biological explanation, and engender in her a reflective distance from the behaviours and beliefs our epistemic discourse prescribes, she may as a result form fewer true beliefs. Though she will have the consolation of having true beliefs about the nature of epistemic normativity.
15. Thanks to Paul Noordhof and Carolyn Price for bringing this way of developing the account to my attention.

Notes on contributor

Emma Sullivan-Bissett is a Lecturer in Philosophy at the University of Birmingham. Her research concerns the nature of belief and its connection to truth, as well as delusional beliefs and how they are formed.

References

- Bar-On, Dorit, and James Sias. 2013. "Varieties of Expressivism." *Philosophy Compass* 8 (8): 699–713.
- Buller, David. 1998. "Etiological Theories of Function: A Geographical Survey." *Biology and Philosophy* 13 (4): 505–527.
- Bykvist, Krister, and Anandi Hattiangadi. 2013. "Belief, Truth, and Blindspots." In *The Aim of Belief*, edited by Timothy Chan, 100–122. Oxford: Oxford University Press.
- Clarke-Doane, Justin. 2012. "Morality and Mathematics: The Evolutionary Challenge." *Ethics* 122 (2): 313–340.
- Côté-Bouchard, Charles. 2016. "Can the Aim of Belief Ground Epistemic Normativity?" *Philosophical Studies*. doi:10.1007/s11098-016-0657-8.
- Cowie, Christopher. 2014. "In Defence of Instrumentalism about Epistemic Normativity." *Synthese* 191: 4003–4017.
- Cuneo, Terence. 2007. *The Normative Web*. Oxford: Oxford University Press.
- Daly, Chris, and Davod Liggins. 2010. "In Defence of Error Theory." *Philosophical Studies* 149: 209–230.
- Engel, Pascal. 2007. "Belief and Normativity." *Disputatio* 2 (23): 179–204.
- Fales, Evan. 1996. "Plantinga's Case Against Naturalistic Epistemology." *Philosophy of Science* 63 (3): 432–451.
- Feldman, Richard. 1988. "Epistemic Obligations." *Philosophical Perspectives* 2: 235–256.
- Joyce, Richard. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.
- Kant, Immanuel. 1785/2005. *Groundwork for the Metaphysics of Morals*. Translated Thomas K. Abbott and edited by Lara Denis. New York: Palgrave Macmillan.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66 (3): 612–640.
- Leiter, Brian. 2015. "Normativity for Naturalists." *Philosophical Issues* 25: 64–79.
- Mackie, John L. 1977. *Ethics: Inventing Right and Wrong*. New York: Penguin Books.
- McKay, Ryan, and Daniel Dennett. 2009. "The Evolution of Misbelief." *Behavioral and Brain Sciences* 32: 493–561.
- Millikan, Ruth Garrett. 1984. *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- Millikan, Ruth Garrett. 1989. "In Defense of Proper Functions." *Philosophy of Science* 56 (2): 288–302.
- Millikan, Ruth Garrett. 1993a. "Explanation in Biopsychology." In *White Queen Psychology and Other Essays for Alice*, edited by R. Millikan, 171–192. Cambridge, MA: MIT Press.

- Millikan, Ruth Garrett. 1993b. "Naturalist Reflections on Knowledge." In *White Queen Psychology and Other Essays for Alice*, edited by R. Millikan, 241–264. Cambridge, MA: MIT Press.
- Millikan, Ruth Garrett. 1993c. "Propensities, Exaptations, and the Brain." In *White Queen Psychology*, edited by R. Millikan, 31–50. Cambridge, MA: MIT Press.
- Neander, Karen. 1991a. "Functions as Selected Effects: The Conceptual Analyst's Defense." *Philosophy of Science* 58 (2): 168–184.
- Neander, Karen. 1991b. "The Teleological Notion of 'Function'." *Australasian Journal of Philosophy* 69 (4): 454–468.
- Nolfi, Kate. 2014. "Why Is Epistemic Evaluation Prescriptive?" *Inquiry* 57 (1): 97–121.
- O'Brien, Lucy. 2005. "Imagination and the Motivational View of Belief." *Analysis* 65 (1): 55–62.
- Olson, Jonas. 2011. "Error Theory and Reasons for Belief." In *Reasons for Belief*, edited by Andrew Reisner and Asbjørn Steglich-Petersen, 75–93. Cambridge: Cambridge University Press.
- Owens, David. 2003. "Does Belief Have an Aim?" *Philosophical Studies* 115: 283–305.
- Papineau, David. 1987. *Reality and Representation*. Oxford: Basil Blackwell Limited.
- Papineau, David. 1993. *Philosophical Naturalism*. Cambridge, MA: Blackwell.
- Papineau, David. 1999. "Normativity and Judgement." *Aristotelian Society Supplementary Volume* 73 (1): 17–43.
- Papineau, David. 2013. "There Are No Norms of Belief." In *The Aim of Belief*, edited by Timothy Chan, 64–79. New York: Oxford University Press.
- Price, Carolyn. 1995. "Functional Explanations and Natural Norms." *Ratio* 8 (2): 143–160.
- Price, Carolyn. 2001. *Functions in Mind*. Oxford: Oxford University Press.
- Quine, W. V. O. 1985. "Natural Kinds." In *Naturalizing Epistemology*, edited by Hilary Kornblith, 57–76. Cambridge, MA: MIT Press.
- Ruse, Michael. 1986. "Evolutionary Ethics: A Phoenix Arisen." *Zygon* 21 (1): 95–112.
- Ruse, Michael. 1993. "The Significance of Evolution." In *A Companion to Ethics*, edited by Peter Singer, 500–510. Oxford: Blackwell.
- Stephens, Christopher L. 2001. "When Is It Selectively Advantageous to Have True Beliefs?" *Philosophical Studies* 105 (2): 161–189.
- Stich, Stephen. 1990. *The Fragmentation of Reason*. London: MIT.
- Street, Sharon. 2009. "Evolution and the Normativity of Epistemic Reasons." *Canadian Journal of Philosophy* 39: 213–248.
- Sullivan-Bissett, Ema. 2016. "Malfunction Defended." *Synthese*. doi:10.1007/s11229-106-1062-8.
- Sullivan-Bissett, Ema. 2017. "Explaining Doxastic Transparency: Aim, Norm, or Function?" *Synthese*. doi:10.1007/s11229-017-1377-0.
- Taylor, S. E., and J. D. Brown. 1994. "Positive Illusions and Well-being Revisited: Separating Fact from Fiction." *Psychological Bulletin* 116 (1): 21–27.
- Trivers, Robert. 2000. "The Elements of a Scientific Theory of Self-deception." *Annals of the New York Academy of Sciences* 907: 114–131.
- Trivers, Robert. 2011. *Deceit and Self-deception: Fooling Yourself the Better to Fool Others*. London: The Penguin Group.
- Trivers, Robert. 2013. *The Folly of Fools. The Logic of Deceit and Self-deception in Human Life*. New York: Basic Books.
- Van Leeuwen, Neil. 2007. "The Spandrels of Self-deception." *Philosophical Psychology* 20 (3): 329–348.
- Van Leeuwen, Neil. 2008. "Finite Rational Self-deceivers." *Philosophical Studies* 139 (2): 191–208.
- Williams, Bernard. 1973. "Deciding to Believe." In *Problems of the Self*, edited by Williams, B. 136–151. London: Cambridge University Press.
- Wrenn, Chase B. 2004. "Hypothetical and Categorical Epistemic Normativity." *The Southern Journal of Philosophy* 42: 273–290.