



Epistemically flawless false beliefs

Kate Nolfi¹ 

Received: 5 May 2020 / Accepted: 7 July 2020 / Published online: 16 July 2020
© Springer Nature B.V. 2020

Abstract

A starting point for the sort of alethic epistemological approach that dominates both historical and contemporary western philosophy is that epistemic norms, standards, or ideals are to be characterized by appeal to some kind of substantively normative relationship between belief and truth. Accordingly, the alethic epistemologist maintains that false beliefs are necessarily defective, imperfect, or flawed, at least from the epistemic perspective. In this paper, I develop an action-oriented alternative to the alethic approach, an alternative that is inspired by and jives with the kind of thinking that underwrites promising and increasingly popular enactive or embodied research programs in cognitive science. Moreover, I argue that the proponent of an action-oriented epistemological approach ought to deny that falsity, in and of itself, necessarily constitutes a kind of epistemic imperfection in belief. The action-oriented epistemologist ought to embrace the possibility that there are epistemically flawless false beliefs.

Keywords Epistemic Evaluation · Epistemic norms · Action-Oriented Epistemology · Alethic · Epistemology

1 Introduction

The sort of alethic epistemological approach that dominates both historical and contemporary western philosophy maintains that truth plays an important and ineliminable role in structuring the domain of epistemic evaluation. In particular, the alethic epistemologist claims, the content of epistemic norms, standards, or ideals is fixed (in some way or other) by appeal to truth as a fundamental epistemic value, goal, or aim. So the alethic approach takes the simple idea that epistemic success is most fundamentally about believing truths and not believing falsehoods as its starting point. Accordingly, the alethic epistemologist maintains that truth is always and inevitably an epistemically good-making feature of belief. That a belief is true necessarily makes the belief

✉ Kate Nolfi
knolfi@gmail.com

¹ The University of Vermont, Burlington, USA

epistemically *better* than it would otherwise be (at least in one respect or along one dimension of epistemic evaluation). In contrast, a belief's falsity inevitably constitutes a kind of epistemic defect, flaw, or failure for the alethic epistemologist. At least from the epistemic perspective, falsity in belief is always a kind of imperfection, and necessarily so.

My goal here is to motivate and explore the implications of an alternative to the alethic approach. I call this alternative the *action-oriented epistemological approach*.

The action-oriented approach takes a familiar and commonsensical idea about what beliefs are as its starting point. This idea, famously expressed by Ramsey's claim that "beliefs are the maps by which we steer," can be expressed in three steps.¹ First, belief has an aim, goal, or proper function. Second, having this aim, goal, or proper function is part of what makes belief the distinctive kind of cognitive state that it is. And third, belief's aim, goal, or proper function must be characterized by appeal to the distinctive ("map-like") way in which beliefs paradigmatically enable or facilitate action.

The kind of action-oriented thinking about the human mind that underwrites this familiar and commonsensical idea about what beliefs are also underwrites the kind of action-oriented framework that stands behind highly productive embodied and enactive research programs in cognitive science. So, if we think that epistemological theorizing ought to jive with the kind of philosophy of mind that guides successful research programs in cognitive science, then we ought to regard the action-oriented epistemological approach as an especially attractive alternative to the traditional alethic approach. Accordingly, I suggest that an action-oriented alternative to the alethic epistemological approach both complements and garners support from the fruitfulness of an action-oriented approach to the empirical study of human cognition. And for this reason, the action-oriented approach merits a kind of serious and sustained philosophical attention that it has yet to receive.

For the action-oriented epistemologist, the most fundamental dimension of epistemic evaluation is not truth (or accuracy), but rather *well-suitedness* to fulfill belief's distinctive action-oriented proper function. If a belief's falsity always and inevitably prevented or precluded the belief from being well-suited to fulfill its action-oriented proper function, then an action-oriented epistemology would entail that a belief's falsity always and inevitably constitutes a kind of epistemic defect, fault, or flaw. But, under certain circumstances and in certain domains, there are reasons to think that some false beliefs are especially—even uniquely—well-suited to fulfill belief's action-oriented proper function. (Or, at the very least, these false beliefs seem to be much better-suited than true beliefs about the same subject matter would be to do the relevant bit of work in the believer's mental economy.) And if this is right, then adopting an action-oriented approach requires us to rethink the epistemic status of false belief. More precisely, I argue that preliminary evidence suggests the action-oriented epistemologist should deny that a belief's falsity, in and of itself, inevitably constitutes a kind of epistemic imperfection. She ought to embrace the possibility that there are epistemically flawless false beliefs: false beliefs that, at least from the epistemic perspective, are entirely without fault or defect.

¹ Ramsey (1927).

Needless to say, this result marks an obvious and seemingly radical point of departure from mainstream alethic epistemology. Moreover, it invites concern that the action-oriented approach amounts to a kind of thoroughgoing pragmatism about cognitive evaluation, one that effectively dissolves the epistemic domain. One might worry that, in characterizing evaluation with respect to standards that derive from belief's constitutive, action-oriented proper function, the action-oriented epistemologist has simply changed the subject; that by relinquishing a commitment to truth as a (perhaps *the*) fundamental epistemic aim, goal, value, or ideal, the action-oriented approach has simply left *epistemology* behind. I show that this cluster of concerns is overblown by arguing that the kind of action-oriented approach that I develop here can respect the autonomy and independence of the epistemic perspective even while allowing that certain false belief may be epistemically flawless.² If my arguments are successful, then the empirically-minded epistemologist has good reason to embrace the kind of departure from epistemological orthodoxy that may be involved in accepting an action-oriented approach, and she can do so without apology.

2 The alethic approach and false belief

The alethic epistemologist's way of thinking is, at least in slogan form, exceedingly familiar: truth (and the avoidance of error) constitutes a kind of fundamental epistemic norm, end, goal, value, or ideal. For the alethic epistemologist, then, truth (and the avoidance of error) both circumscribes and determines the internal structure of the epistemic domain.³ Many alethic epistemologists treat the idea that truth constitutes a kind of fundamental epistemic norm, end, goal, value, or ideal as a kind of conceptual or analytic fact. For these theorists, that truth plays the relevant sort of fundamental normative role in circumscribing and structuring a distinctive evaluative perspective is precisely what distinguishes the relevant perspective as *epistemic* (rather than moral, aesthetic, financial, pragmatic, etc.) in the first place. Of course, one might simply treat this claim about the role truth plays in circumscribing and structuring the epistemic domain as the expression of a bottom-level analytic or conceptual truth—the kind of truth that neither requires nor admits of further explanation. But if, as I suspect is the case, our ordinary, everyday ways of evaluating believers, beliefs, and belief-regulation do not clearly or univocally appeal to truth as the fundamental norm, goal, aim, value, or ideal for belief, then we have some reason to resist treating as basic (i.e. as unexplained and unexplainable) that an epistemic perspective qualifies as *epistemic* by virtue of the normative role that truth plays in structuring and circumscribing this

² As I suggest below, I think we should resist the idea that the epistemic perspective just is, as a matter of definition or basic conceptual/analytic fact, the evaluative perspective within which truth serves as a/the fundamental aim, goal, value, or ideal. But for the reader who is unwilling to relinquish this commitment, the upshot of this paper's discussion is that (i) an action-oriented approach exposes a coherent alternative evaluative perspective, and (ii) this alternative evaluative perspective is especially well-positioned to make sense of our ordinary, everyday, and pretheoretical doxastic evaluative practices and, as such, merits philosophical attention.

³ The alethic approach, although popular, is not altogether uncontroversial in contemporary epistemology. Perhaps the most salient alternative is the kind of knowledge-first approach associated with the work of Timothy Williamson, among others.

perspective. In any case, it is reasonable to hope for epistemological theorizing to do more: the epistemologist should hope to be able to explain *why* truth plays whatever role it plays in circumscribing and structuring the epistemic domain. Put differently, the epistemologist should hope to be able to give a non-trivial answer to questions like: “Why does truth have the kind of normative authority in the epistemic domain that it has?” and “Why should truth serve as a fundamental norm, goal, aim, ideal, or value in our ordinary epistemic practices evaluating believers, beliefs, and belief-regulation?”⁴

One popular strategy that alethic epistemologists have pursued in an effort to explain why truth plays the relevant sort of fundamental normative role in circumscribing and structuring the epistemic domain unfolds as follows. The alethic epistemologist begins by embracing an intuitively attractive and widely-endorsed methodological commitment:

MC: we ought to look to an account of what beliefs *are* to explain the content of epistemic norms or standards, and so to explain the structure and character of epistemic evaluation.⁵

Then, and rather predictably, the alethic epistemologist couples *MC* to a familiar way of thinking about the (metaphysical and/or conceptual) nature of belief. In slogan form: belief *constitutively aims* at truth.⁶ A bit more carefully:

The Alethic Account of Belief: belief, by virtue of being the kind of cognitive state that it is, aims (where this notion of “aiming” might be construed in any number of different ways) at truth.

An alethic epistemology treats the fact that belief constitutively aims at truth as explanatorily fundamental or brute. And it derives other epistemic norms, requirements, ideals, or standards from belief’s truth aim. The defining commitment of the alethic approach, then, is that the content of the epistemic norms, requirements, ideals, or standards, and so the structure and character of epistemic evaluation, must (ultimately) be grounded in, and so explained by appeal to, the fact that belief aims at

⁴ If one takes on board the assumption that the epistemic domain just is, as a matter of basic conceptual or analytic fact, the domain for which truth serves as the normatively fundamental structuring goal or value, then the epistemologist ought to be able to explain why theoretical work aimed at revealing the internal structure of an evaluative domain for which truth serves as the normatively fundamental goal, aim, or value helps us understand and/or improve our ordinary, everyday evaluations of believers, beliefs, and belief-regulation. Put differently, the epistemologist ought to be able to explain why and how the epistemic perspective (so conceived) ought to figure in our ordinary, everyday evaluative practices. And so the epistemologist still owes an explanation of why truth should serve as a fundamental norm, goal, aim, ideal, or value in our ordinary, everyday evaluations of believers, beliefs, and belief-regulation.

⁵ Although quite popular (see, e.g., those authors cited in note 6), this methodological thesis, associated with what we might call a doxastic account of the epistemic, is not wholly uncontroversial (see Hazlett (2013) for useful discussion). I won’t try to motivate, much less defend, *MC* in what follows. My more conservative ambition is to show that if one accepts this methodological thesis, as many epistemologists seem to, then one has good reason to embrace the kind action-oriented approach to epistemological theorizing over the currently and historically popular alethic approach.

⁶ There is no consensus among those who adopt this kind of alethic approach regarding how best to characterize this norm. Perhaps a belief *ought* to be true. Alternatively, perhaps a belief is *correct* or *successful* (if and) only if true. Or perhaps S ought to believe that *p* (if and) only if it is true that *p*. (See, e.g. Gibbard (2005), Shah (2003), Shah and Velleman (2005), Sosa (2009), Velleman (2000), and Wedgwood (2002, 2007, 2013), among others).

truth. Moreover, for the alethic epistemologist, it is the fact that certain norms (e.g. norms of epistemic rationality or justification) are grounded in or entailed by belief's constitutive truth aim that makes these norms genuinely *epistemic*.

Any version of the alethic approach straightforwardly entails that false belief inevitably and necessarily falls short along one (particularly fundamental) dimension of epistemic evaluation. After all, if belief constitutively aims at truth, then a false belief is always and inevitably epistemically imperfect, defective, or flawed by virtue of its failure to achieve its constitutive aim. Thus, the alethic epistemologist is committed to the claim that a false belief is always appropriately subject to at least one fundamental species of epistemic criticism or negative evaluation, and necessarily so.

3 Motivating an action-oriented alternative

My ambition here is to articulate and motivate an underexplored alternative to the alethic approach.⁷ This alternative is inspired by a line of thought that has roots in the American pragmatist tradition and serves as the theoretical foundation for an impressive body of research in contemporary cognitive science. Accordingly, let me begin by simply trying to isolate and precisify the relevant inspirational line of thought from the contemporary literature in cognitive science.

Andreas Engel et al. articulate what they take to be the defining core commitment of a “pragmatic turn” in cognitive science:

...cognition should not be understood as [for the purpose of] providing models of the world, but as subserving action... Accordingly, cognitive processes and their underlying neural activity patterns should be studied primarily with respect to their role in action generation.⁸

And across the literature, cognitive scientists express this same core commitment as follows:

[M]inds evolved to make things happen...the biological mind is, first and foremost, an organ for controlling the biological body...[m]inds are not disembodied logical reasoning devices.⁹

[T]he primary role of cognition is to support (and enhance) action-control abilities rather than to produce ‘encyclopedic’ knowledge that is detached from action and perception systems. All cognitive operations...are organized around—and

⁷ There are at least two other alternatives to alethic epistemology on offer in the literature that bear mentioning here: the so-called knowledge-first approach (see, e.g., Williamson (2000)) and what we might call the rationality-first approach (see, e.g., Gibbons (2013)). I simply do not have space to give either of these alternatives the careful treatment each deserves. But note that if the knowledge-first epistemologist accepts that knowledge is factive, then she will share the alethic epistemologist's commitment to the claim that false beliefs are always and inevitably epistemically faulty. And so if, as I argue below, empirical results pressure the action-oriented epistemologist to deny this, then it seems that the action-oriented approach stands in tension with the most prominent version of this alternative, not only at the level of epistemological theory, but also with respect to the verdicts each issues regarding the epistemic status of particular beliefs.

⁸ Engel et al. (2013).

⁹ Clark (1997).

ultimately functional to—the demands of action control and goal achievement...¹⁰

The function of the mind is to guide action, and cognitive mechanisms... must be understood in terms of their ultimate contribution to situation-appropriate behavior.^{11,12}

If, as these quotations suggest, cognition proceeds, in the first instance or most fundamentally, in the service of action, then we should expect the architecture of the cognitive system to subserve action-production in the same way that the design and construction of a household flashlight subserves the flashlight's particular design goal: portable production and projection of light. Both the structure (e.g. the wiring, the material and shape of the parabolic reflector behind the bulb) and the proper operation (e.g. the way in which switching the flashlight on ought to complete an electrical circuit and thereby illuminate the filament in the bulb) of the flashlight are shaped by, and so most fruitfully understood in terms of, the peculiar demands of the flashlight's design goal. Similarly, then, if the purpose of cognition is to facilitate successful action, then understanding the human cognitive system requires conceiving of the cognitive system as a system that serves the peculiar demands of action-production in human organisms operating in normal environments.¹³

We can usefully precisify this line of thought by construing it as a commitment to an *action-first* metaphysical thesis about the nature of cognitive state types. The thesis is this: that the way in which a particular type of cognitive state paradigmatically operates within a subject's mental economy to produce or sustain action determines the distinctive nature of that cognitive state type—it is that in virtue of which the relevant type of cognitive state is the distinctive type that it is. If this thesis is correct, then cognitive state types, as well as the cognitive process types that generate and sustain these states, can be individuated by appeal to, and must be characterized in terms of, the contributions that they paradigmatically make in generating and sustaining action.

In functionalist terms, any particular type of cognitive state has a distinctive proper function. And this proper function partially constitutes (metaphysically speaking) the relevant type as the type that it is. Moreover, this proper function is action-oriented in the sense that this proper function, and so the proper functioning of the particular cognitive processes that generate or sustain states of the relevant type, are determined by, and so must be characterized by appeal to, the particular, distinctive contribution that states of the relevant type paradigmatically make to action-production.

This action-first metaphysical thesis entails that the nature of belief, as well as the character of the cognitive processes that regulate belief, can only be characterized by appeal to what it takes for belief to be well-suited to play the distinctive role

¹⁰ Pezzulo (2015).

¹¹ Wilson (2002).

¹² For further discussion of this core commitment as a theoretical foundation for empirical research in cognitive science, see, e.g., Engel, Friston, and Kragic (eds.) (2015).

¹³ A growing number of cognitive scientists treat this line of thought as a point of departure for their empirical work, and to impressive effect. For a review of the wealth of empirical work that is inspired by and grounded on this action-oriented line of thought, see, e.g., Shapiro (2010).

(whatever it is) that beliefs paradigmatically play in action-production. Call this the *Action-Oriented Account of Belief*.

I think the Action-Oriented Account of Belief points us toward an action-oriented alternative to the alethic approach in epistemology. And, to the extent that we think epistemological theorizing ought to jive with the kind of thinking that underwrites fruitful research programs in cognitive science, the action-oriented account of belief motivates this alternative.¹⁴ This action-oriented alternative shares the alethic approach's commitment to *MC*, and so the action-oriented epistemologist aims to ground her account of the structure and character of epistemic evaluation in an account of what beliefs *are*. But instead of adopting the Alethic Account of Belief, the action-oriented approach couples *MC* to the Action-Oriented Account of Belief. And, accordingly, the action-oriented approach takes the fact that beliefs constitutively aim, in some to-be-precisified sense, at being well suited to subserve action in a particular, distinctive way as the point of departure for epistemological theorizing.¹⁵

Notice that the action-oriented epistemologist can easily identify and describe a system of constitutive norms which govern belief and belief-regulation. Most fundamentally, on the action-oriented approach, the thesis that beliefs have the particular, action-oriented proper function that they do entails that a belief enjoys a kind of positive epistemic status—a belief is appropriately subject to a species of epistemic praise; there is something good about it, epistemically speaking—when it is well-suited to fulfill this action-oriented proper function. And any belief that is not well-suited to fulfill belief's proper function is faulty or defective *qua* belief (and so thereby criticizable along at least one salient dimension of epistemic evaluation). The action-oriented epistemologist might capture this thought by saying that belief, simply by virtue of being the kind of cognitive state that it is, is governed by a certain explanatorily fundamental norm of correctness: a belief is correct to the extent that it is well-suited to fulfill belief's action-oriented proper function. Additionally, belief's constitutive action-oriented proper function underwrites further constitutive norms governing belief-regulation. These are norms of ideal cognitive functioning with respect to belief-regulation for believers equipped with our cognitive equipment and our cognitive limitations, operating in an environment like our own, and whose mental economies are wired up in the service of action-production in the way that ours are.¹⁶ They characterize the particular ways in which cognition ought to regulate our beliefs—they map those ways in which cogni-

¹⁴ I've presented other, independent arguments in defense of this sort of action-oriented alternative in Nolfi (2015, 2019).

¹⁵ Proponents of this sort of action-oriented approach include Burge (2010), Kornblith (2002, 1993), Lycan (1988), Millikan (1993), Nolfi (2015, 2018a, b, 2019), and Papineau (1987).

¹⁶ By way of illustration: a consequentialist version of the action-oriented approach might maintain that norms of epistemic rationality and justification are just norms that specify patterns of belief regulation (i.e. patterns of transition from inputs to cognition—perceptual experiences, other beliefs, etc.—to cognitive output—namely, belief) instantiations of which best, most effectively, or most reliably equip believers like us with beliefs that are well-suited to fulfill their proper function. But the action-oriented epistemologist need not embrace this kind of epistemic consequentialism—it is open to the action-oriented epistemologist to embrace a non-consequentialist account of the way in which derivative epistemic norms are grounded in the fact that beliefs are praiseworthy along the most fundamental dimension of epistemic appraisal when they are well-suited to fulfill belief's action-oriented proper function (against epistemic consequentialism, see, e.g., Berker (2013)).

tive functioning most effectively translates any particular set of inputs (e.g. perceptual experience that p) to doxastic output (e.g. belief that p)—so as to equip us with beliefs that are well-suited to play the distinctive role that beliefs paradigmatically play in the course of action-production. And crucially, on the action-oriented approach, these norms—i.e. the norms that govern belief in virtue of belief's having the particular metaphysical nature that it does—are just the norms evaluation with respect to which is genuinely *epistemic* in character. On the action-oriented approach, then, epistemic norms governing belief-regulation specify how ideal cognition ought to proceed in creatures who have the kind of cognitive equipment that we have, and who operate in the environment in which we operate, given the distinctive contribution that beliefs paradigmatically make to action-production in mental economies that are structured as ours are.

4 Developing the action-oriented approach

For the proponent of the action-oriented approach, whether or not false belief is always and inevitably epistemically defective or flawed depends on the nature of the distinctive role that beliefs paradigmatically play in the course of action-production, and so on the particular details of the account of belief's distinctive proper function that the action-oriented epistemologist accepts. Of course, determining the proper function of belief is, at least in part, an empirical endeavor. So, developing an action-oriented epistemology that delivers verdicts on the epistemic status of false beliefs requires empirical inquiry and is hostage to empirical results.

That said, I think we already have compelling evidence that a certain empirical hypothesis about the role beliefs paradigmatically play in the course of action-production is at least roughly correct. This hypothesis states that beliefs inform our actions by allowing us to anticipate the outcomes of various potential actions across a wide variety of different situations so as to facilitate action selection that results in the achievement of our different goals. Our beliefs put us in a position to predict the way our dynamic environment will evolve over time (and especially as a result of the various courses of action we might perform). And these predictions make it possible for us to select actions that will efficiently and effectively achieve our various ends given our particular circumstances. In slogan form, belief is a kind of flexible-use predictive tool for action selection; a tool that can be deployed across a wide variety of contexts and in the service of a variety of different ends to generate predictions about the outcomes of potential actions so as to engender the agent's selection of an actions that achieves the agent's particular ends at the time of action.¹⁷

¹⁷ One natural way to frame this hypothesis appeals to Ramsey's metaphor: paradigmatically, at least, beliefs inform action in a similar way to the way in which a subway map informs a commuter's efforts to navigate a city. The commuter relies on the map in order to predict, e.g., where she will end up if she takes the D train to the end of the line, whether she can expect to find a D train on the platform at the 34th Street station, or which stops the D train will make between West 4th and Columbus Circle. And, regardless of where she finds herself in the city, the predictive power of the subway map will allow the commuter to plan her subway ride so that she gets wherever she needs to go (assuming the world cooperates, of course). Similarly, our beliefs equip us with a kind of flexibly deployed predictive power. Regardless of where we happen to find ourselves, this predictive power puts our cognitive systems in a position to identify courses of

Recall that, on the action-oriented approach, our beliefs are epistemically correct insofar as and to the extent that they are well-suited to fulfill belief's action-oriented proper function. And norms of ideal cognitive functioning with respect to belief regulation are just those norms conformity with which best equips believers who have the kind of cognitive equipment that we have and operate in the kind of environment in which we operate with beliefs that are well-suited to fulfill belief's particular action-oriented proper function. So, if the flexible-use predictive tool hypothesis is at least roughly correct, then the action-oriented approach entails that a belief is epistemically correct or successful insofar as and to the extent that it is well-suited to inform our actions by serving as the relevant sort of predictive tool. And norms of ideal cognitive functioning with respect to belief regulation are just norms that capture patterns of cognitive processing which most reliably yield (as output) beliefs that are epistemically correct or successful in just this sense.

The beliefs that will be most well-suited to serve as the relevant sort of flexible-use predictive tool in action-production will be poised to facilitate successful action across the wide range of circumstances in which human beings normally find themselves, and in the service of a wide variety of ends that human beings normally aim to achieve. In short, they will be versatile or multipurpose. But, their versatility must be bounded. After all, if a tool—any tool—is to be *useful*, it must be calibrated to the capabilities of its user, to a circumscribed set of tasks the tool is meant to facilitate, and to a restricted set of circumstances in which the tool is to be put to use. A household flashlight can only do its particular job well if it is easy for an average human being to grasp and carry, if its bulb produces light in the visible spectrum (as opposed to ultraviolet or infrared light). The more precisely a tool is attuned or tailored to the capabilities and limitations of its target user, to a circumscribed set of tasks, and to a restricted set of circumstances in which the tool is to be put to use, the less versatile the tool will be.

So, if human beliefs are flexible-use predictive tools, then that ideal human belief is belief which is optimized to facilitate prediction in the service of action selection in the sorts of circumstances in which human beings paradigmatically find themselves, for the sorts of ends that human beings paradigmatically have, and for what it typically takes for creatures who are wired up in the way that human beings are to successfully pursue these sorts of ends. Accordingly, epistemic evaluation will prioritize or privilege being well-suited to play the relevant predictive role in action-production within mental economies that have the distinctive structure and character that ours have, in the sorts of circumstances in which human beings paradigmatically find themselves, and in pursuit of the sorts of ends that human beings paradigmatically pursue.¹⁸

Footnote 17 continued

action that will efficiently and effectively achieve our various ends, whatever those ends may be. Although this metaphor has virtues, the analogy it suggests is not perfect. After all, we do not always, or even typically, consult our beliefs consciously, effortfully, and explicitly in the course of coming to act in the way that the (novice) commuter consults the subway map as she plans her trip. Additionally, the representational content of a subway map appears to be (at least in part) non-propositional, but beliefs certainly can (and perhaps must) have propositional content.

¹⁸ This discussion brings out the action-oriented epistemologist's commitment to the thesis that epistemic norms are norms which govern (at least in the first instance) human belief and human believers in normal worlds.

5 Action-oriented epistemology and false belief

What does it take, then, for human beliefs to be well-suited to serve as flexible-use predictive tools in action-production in the sorts of circumstances in which human beings paradigmatically find themselves, and in pursuit of the sorts of ends that human beings paradigmatically pursue, given the character and structure of human mental economies?¹⁹

At first pass, it might appear obvious that only true beliefs could be genuinely well-suited to serve as flexible-use predictive tools in action-production.²⁰ A subway map that misrepresents the order of the stops on the subway lines will be notably unhelpful when it comes to navigating the relevant portion of the city. And the reason is precisely that predictions (e.g. about where one will be if one gets on the D train heading uptown at West 4th and rides for three stops) that such a map puts us in a position to make are poorly-suited to guide a traveler's efforts to get herself from one location to another via the subway. One way of developing the action-oriented approach seizes on and generalizes from this sort of insight in an effort to vindicate the claim that belief, indeed, aims at truth (albeit derivatively so).²¹ This way of developing the action-oriented approach rests on the hypothesis that a belief will only be well-suited to fulfill its action-oriented proper function if it is true. And if this hypothesis is right, then the fact that belief is governed by an fundamental correctness-as-well-suitedness-to-fulfill-beliefs'-proper-function epistemic norm can explain why belief is governed by a epistemic norm demanding truth.²² Any belief that is well-suited to fulfill belief's particular action-oriented proper function will have to be true. A belief's truth, then, is a prerequisite for that belief's being well suited to fulfill its action-oriented proper function. Accordingly, on this way of developing the action-oriented approach, an action-oriented epistemology and an alethic epistemology converge: if a belief's being well-suited to fulfill its action-oriented proper function requires that the belief be true, then the action-oriented epistemologist will agree with her alethic counterpart that false beliefs are guaranteed to fall short with respect to at least one

¹⁹ This is, of course, an (at least in part) empirical question, and so adopting the action-oriented approach involves embracing the idea that what it takes to conform with epistemic norms is to be discovered (at least in part) through empirical inquiry. Thus, for the action-oriented epistemologist, the possibility of epistemically flawless false beliefs depends on (indeed, is hostage to) empirical results. To be clear: the relevant empirical facts are far from settled. Nevertheless, I argue here that an initial survey of the relevant empirical data leaves the action-oriented epistemologist with good reason to be skeptical that false belief is inevitably and necessarily epistemically flawed.

²⁰ After all, as Quine (1969) puts it, "creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind".

²¹ Kornblith (2002, 1993), Lycan (1988), Millikan (1993), and Papineau (1987) develop the action-oriented approach along precisely these lines.

²² In fact, one might be tempted to think there is an even tighter connection between accurate representation and well-suitedness to fulfill beliefs' proper function. Indeed, it can seem that accurate representation and well-suitedness to fulfill beliefs' action-oriented proper function amount to the same thing. And if this is so, then the action-oriented proper function of belief *is just* to accurately represent the facts. I suspect that a (sometimes suppressed—see Millikan 1993, sometimes explicit—see Kornblith 2002) assumption that accurate representation *is just what it takes* for beliefs to be well-suited to serve as predictive tools in guiding action underwrites the popular thought that beliefs constitutively aim to accurately represent the facts. But, as I argue below, this assumption runs counter to empirical evidence.

(particularly fundamental) dimension of epistemic evaluation. And thus she will accept that false beliefs must always and inevitably be epistemically faulty or flawed.

But there is good reason to doubt that a belief's being well-suited to fulfill its action-oriented proper function requires that the belief be true. (Or at least I will try to demonstrate as much in what follows.) Accordingly, I suggest we have reason to be skeptical that an action-oriented epistemology will yield the result that false beliefs are guaranteed to be epistemically faulty or flawed simply by virtue of their falsity.

Although far from conclusive, empirical research in psychology and cognitive science on positive illusions supply one motivation for skepticism. Widely accepted interpretations of the results of this body of research suggest that we may be, as a general rule, more successful in achieving our various ends when our beliefs about ourselves and about our ability to control the world around us are mildly, but systematically, distorted in specific ways. For instance, if my belief corpus includes slightly optimistically inflated representations of the degree to which I am (as well as the degree to which others think of me as) intelligent, hard-working, resilient in the face of adversity, kind, caring etc.—if, that is, my beliefs about myself and about the way in which others view me generally code a mild, but systematic distortion of the relevant facts—then it seems that I may, e.g., have an easier time making friends, perform better in job interviews, and be more successful in efforts to convince others to invest in my business venture as a result. And if my belief corpus slightly overestimates the degree to which I am responsible for the good things that happen to me and slightly underestimates the degree to which I am responsible for the bad, then it seems that I may, as a result, be more resilient in the face of tragedy and hardship. Thus, it seems that I may be more successful in achieving various of my ends across a wide range of different circumstances if my beliefs in certain domain(s) systematically distort the facts in specific ways than I would be were my beliefs in these domain(s) to accurately represent the facts.²³

By way of illustration: imagine that S's optimistic self-regarding beliefs lead S to slightly overestimate S's chances of success in some particular endeavor. Of course, if S were to grossly overestimate S's chances of success here—if, in effect, S were to grossly underestimate the difficulty of achieving S's particular goal—then it is unlike that S will succeed in the endeavor at hand. But if S's optimism is only slightly unrealistic—if, that is, the kind of systematic distortion in S's self-regarding beliefs that S exhibit is mild—then perhaps S's optimism better equips S for success than a realistic assessment of S's chances would. Perhaps, for the sake of illustration, S's optimism makes S more likely to commit to the endeavor at hand, persevering in the face of surmountable, but daunting challenges, than S would be were S to accurately predict S's chances of success. Of course, S's optimism leads S to make predictions about the likely outcomes of various courses of action S might perform which are, strictly speaking, false. But, within S's mental economy, the falsity of these predictions would effectively counteract the effects of a general aversion to risk and a fear of failure on the psychological mechanisms that control action selection. As a result, S selects a course of action that is more likely to result in the achievement

²³ For an overview of the relevant psychological results, see, e.g., Taylor & Brown (1988, 1994), Johnson and Fowler (2011), Sharot (2011), and Bortolotti and Antrobus (2015). See Hazlett (2013) or McKay & Dennett (2009) for a philosophical discussion of some of the relevant psychological research.

of S's current goal than is the course of action S would have selected had S more accurately estimated her own chances of success. And so it is a consequence of S's distinctive psychological constitution that certain of S's strictly-speaking false self-regarding beliefs are particularly well-suited to serve as predictive tools in the course of action selection not in spite, but precisely because of the way in which they distort the facts.

Plausible interpretations of the results of empirical work on positive illusions, then, give us reason to suspect that normal adult human beings are, in fact, psychologically constituted such that, at least in some (relatively encapsulated) domains, certain beliefs which are moderately (but not wildly) skewed or distorted furnish the human cognitive system with *better* predictive tools for action-selection than their truthful counterparts could. Indeed, these false beliefs make for better predictive tools precisely because of the ways in which they distort the truth. They are particularly well-suited to fulfill belief's proper function not by virtue of supporting true predictions, but rather by virtue of supporting predictions capable of facilitating successful action selection *within the broader context of human psychological constitution*. Moreover, there is some reason to think that our cognitive systems function properly when they *introduce* and *maintain* the relevant kinds of (again, relatively encapsulated and domain-specific) systematic distortions into the way in which our belief corpus represents the facts. At least for creatures who are psychologically constituted in the way that we are, it seems that sometimes ideal cognitive functioning may well be straightforwardly unreliable in the alethic epistemologist's sense of the term: some ideal belief-regulating cognitive processes—i.e. some forms of ideal cognitive functioning—properly generate and sustain beliefs that are, strictly-speaking, false.²⁴

Empirical work on positive illusions, then, gives the action-oriented epistemologist some reason to be skeptical of the claim that truth is a universal prerequisite for belief to be well-suited to fulfill its particular action-oriented proper function, at least in creatures who are psychologically constituted in the way that human beings are. But philosophical work on the nature of friendship, the nature of faith (in others, as well as in ourselves), and on the nature of diachronic agency may well supply additional reasons for doubt.

A number of contemporary theorists have suggested that these ways of relating to others and to ourselves may constitutively involve adopting a mildly, yet systematically, distorted view of one's friends, those in whom one has faith, or oneself. That is, friendship, faith, and the kind of agential perspective on our own future actions that makes it possible to promise or resolve in the face of temptation, each plausibly require some form of what has come to be known as epistemic (i.e. doxastic) partiality.²⁵

²⁴ See note 23 for references. Bortolotti and Antrobus (2015) provide an admirably careful discussion (on which this discussion relies) of why, how, and to what extent current empirical work lends support to the claim about human psychological constitution under consideration here.

²⁵ On friendship, see Keller (2004) and Stroud (2006). On faith, see Preston-Roedder (2018). On the sort of diachronic agency involved in promising and resolving in the face of temptation without succumbing to bad faith, see Marušić (2015). It is worth noting that these authors do not argue for (and, indeed, some would be decidedly unsympathetic to) the claim that epistemic norms can or should make room for the kinds of doxastic partiality that help to constitute genuine friendship, faith, and diachronic agency. In effect, however, I argue here that this is precisely what the action-oriented approach makes possible, and perhaps even demands.

By way of illustration: Some have suggested that friendship constitutively involves or requires a disposition to give one's friend the benefit of the doubt by interpreting available evidence so as to put the friend's behavior in a favorable light. If my very close friend has been accused of some sort of moral infraction, and the available evidence points overwhelmingly to her guilt, then perhaps our friendship demands that I be slightly less confident of my friend's guilt (and perhaps more confident that she is being framed) than a third party with access to the same body of evidence might reasonably be. The idea is not that my beliefs here ought to manifest insensitivity to the evidence. Quite to the contrary, the idea is that I ought to regulate my beliefs about my friend's guilt in response to the evidence, and the way in which I ought to respond to the evidence here diverges (systematically, but perhaps only by small degrees) from the way in which a neutral party ought to respond to the same body of evidence in regulating their beliefs about my friend's guilt. Our friendship demands that I treat evidence of my friend's guilt as less probative or less weighty in assessing the likelihood of my friend's guilt than a neutral party should.

If it is right that friendship, faith, or the kind of diachronic agency involved in promising or resolving in the face of temptation constitutively involve this sort of doxastic partiality, then friendship, faith, and diachronic agency demand systematically biased (and, strictly speaking inaccurate) belief. But this result entails that certain of the beliefs that best equip us to participate in friendship, to have faith in ourselves and others, and to understand ourselves as the authors of our actions are beliefs that misrepresent reality in mild, yet systematic ways. And so if—as seems plausible—friendship, faith, or understanding ourselves as agents are (or are prerequisites for achieving) paradigmatic human ends, then it seems that achieving some paradigmatic human ends requires having certain beliefs that distort the facts. Accordingly, the action-oriented epistemologist finds additional evidence here which suggests that some beliefs (and some belief-regulating mechanisms) which distort the facts in certain restricted, but systematic ways can better serve as predictive tools in human cognitive economies than their accurate counterparts could.

So, empirical work on human cognition and philosophical work on the nature of friendship, faith, and the kind of diachronic agency involved in promising or resolving in the face of temptation yield some preliminary evidence for (i) the thesis that some beliefs which are, strictly speaking, false may well be especially well-suited to fulfill belief's action-oriented proper function in mental economies like ours, and (ii) the thesis that some belief-regulating mechanisms which are, strictly-speaking, designed to engender and sustain misrepresentation may most effectively equip believers like us with beliefs that are well-suited to fulfill belief's proper function.²⁶ And this suggests that pursuing the action-oriented approach in a way that is appropriately responsive to

²⁶ It may be that any kind of distortion in one's belief corpus, however mild, puts the subject at risk vis-à-vis acting successfully: a distortion always has the potential (in the right circumstances) to frustrate one's ability to realize one's ends. Nevertheless, if the arguments above are successful, then it seems we have good reason to think that (at least in certain domains) holding beliefs that encode mild distortions into one's belief corpus is decidedly *less risky* than holding their accurate counterparts would be. Although there are hypothetical circumstances in which the relevant mildly distorting beliefs play their paradigmatic role in action-production in a way that leads us to act poorly, these hypothetical circumstances are far less likely to be circumstances that we actually encounter than are circumstances in which the relevant mildly distorting beliefs play their paradigmatic role in action-production in a way that leads us to act well.

this evidence will involve embracing the thesis that certain false beliefs may well be epistemically correct, and certain mechanisms of belief-regulation which routinely generate and sustain false beliefs instantiate epistemically ideal belief-regulation. Thus, the action-oriented epistemologist ought to remain open to the possibility that some false beliefs may well be epistemically flawless.

Of course, it is plausible that true beliefs are *often* or *typically* better-suited to fulfill belief's particular predictive role in our mental economies than are beliefs that misrepresent the facts. It is also plausible that holding wildly inaccurate beliefs—i.e. beliefs which significantly distort the facts—will, putting exceptional circumstances to one side, substantially and systematically handicap a one's ability to act successfully in the pursuit of one's ends, whatever these may be. Accordingly, it may very well be that epistemically flawless false beliefs are anomalous.²⁷ But even if these epistemically flawless false beliefs are anomalous and/or rare, the action-oriented epistemologist has reason to deny that false beliefs are always and inevitably epistemically flawed. Moreover, she has reason to reject this thesis both when it is construed as a necessary truth, built into the nature of belief from the get-go, as the alethic epistemologist suggests, and when it is construed as an empirical truth, grounded in what it takes for belief to be well-suited to inform successful action by serving as a kind of flexible-use predictive tool.²⁸

6 Distinguishing action-oriented epistemology from straightforward pragmatism about cognitive evaluation

Because she is committed to understanding what it takes for a belief to be epistemically praiseworthy (or otherwise) in terms of the way in which beliefs paradigmatically facilitate successful action, the action-oriented epistemologist might sensibly be classified as a kind of pragmatist about epistemic evaluation. Nevertheless, an action-oriented epistemology need not and should not accept that any false belief which happens to be well-suited to facilitate a particular subject's successful action in the circumstances she faces is thereby epistemically flawless. And so the action-oriented approach contrasts with what we might call *straightforward pragmatism* about cognitive evaluation: the view that that a S's belief that *p* is "epistemically" virtuous or praiseworthy *only* insofar as and to the extent that it actually facilitates or would facilitate S's successful action

²⁷ The research on positive illusions, for example, suggests that epistemically flawless false beliefs crop up in specific, fairly circumscribed, and relatively encapsulated domains (e.g. self-belief, belief about the causal history or future likelihood of positive and negative events in one's own life and the lives of one's friends). See references in note 23.

²⁸ If believing constitutively involves *taking* or *regarding as* true (and I have no interest in denying as much), this result raises a cluster of questions about how action-oriented epistemic norms can figure in the first-person deliberative perspective. Unfortunately, I simply do not have space to properly address this cluster of questions in this paper. I have made some preliminary efforts along these lines in Nolfi (2018a and b), but more remains to be said.

in the particular circumstances *S* happens to face.^{29,30} The straightforward pragmatist effectively denies that there is an evaluative perspective that is *distinctively epistemic* in character—what we might be tempted to call “epistemic” evaluation, on her view, is *just* (a species of) practical or prudential evaluation. An action-oriented epistemology, however, can avoid this counterintuitive result.

A familiar sort of case will make the contrast between an action-oriented epistemology and straightforward pragmatism about cognitive evaluation vivid. Imagine Bella knows that she can secure substantial benefit or avoid substantial harm merely by believing that 2343 is a prime number, a false proposition, and one in support of which Bella has absolutely no evidence. Bella’s future actions will be significantly more successful if she believes that 2343 is prime than if she fails to so believe.³¹ For this reason, a straightforward pragmatist about cognitive evaluation is committed to saying that, given her circumstances, Bella ought to believe that 2343 is prime. For the pragmatist, the fact that the belief that 2343 is prime is uniquely well-suited to facilitate Bella’s successful action going forward entails that the belief that 2343 is prime is the correct belief for Bella to have. And should Bella manage to form the belief that 2343 is prime (perhaps she pulls off this feat via hypnosis or some other form of cognitive manipulation), it seems the straightforward pragmatist must deny that Bella’s belief is an appropriate target for any kind of negative “epistemic” evaluation.³² Rather the opposite: the straightforward pragmatist should regard Bella’s belief that 2343 is prime as genuinely and entirely praiseworthy, at least from the “epistemic” perspective.

In contrast, the action-oriented epistemologist can vindicate our intuitive response to Bella’s case: even though the belief that 2343 is prime is uniquely positioned, given Bella’s peculiar circumstances, to facilitate Bella’s successful action going forward, Bella would, nevertheless, be subject to epistemic criticism were she to believe that 2343 is prime. An action-oriented epistemologist might ground epistemic criticism of Bella’s useful-but-false belief by pointing out that the belief in question wouldn’t have similarly facilitated successful action across the wide variety of different circumstances human beings normally face, and so is not, in fact, well-suited (in the relevant sense) to fulfill belief’s particular action-enabling function. After all, circumstances in which one finds that one stands to receive a large monetary reward for adopting any particular beliefs about whether a particular four-digit number is prime are certainly abnormal. Alternatively, the action-oriented epistemologist might cite the fact that Bella’s useful-but-false belief is likely generated and sustained by belief-regulating processes that do not reliably give rise to beliefs that are well-suited to fulfill belief’s

²⁹ Views in the vicinity of straightforward pragmatism have been advanced by Stephen Stich (1993), Miriam McCormick (2014) and Susanna Rinard (2015).

³⁰ It is, perhaps, helpful to note that if the action-oriented epistemologist were to adopt a consequentialist approach to developing her account of epistemic norms, then the distinction between what I call straightforward pragmatism and the resulting consequentialist version of action-oriented epistemology would structurally parallel the distinction between act- and rule-consequentialism in ethics.

³¹ Assume, if you like, that Bella has no great interest in large prime numbers, and that Bella is not (and does not aspire to be) a number theorist.

³² Note: the straightforward pragmatist can certainly allow that an “epistemically” flawless belief might criticizable along some other, independent dimensions of evaluation (e.g. such a belief might be aesthetically, or perhaps even morally, objectionable).

action-enabling function as grounds for a kind of negative epistemic evaluation.³³ Plausibly, whatever belief-regulating processes are responsible for Bella's useful-but-false belief manifest a kind of objectionable insensitivity to available evidence (or to a lack thereof). Insofar as she manages to get herself to believe that 2343 is prime, then, Bella gets lucky: she regulates her beliefs in a way that would have gotten her into trouble had her circumstances been more typical. Indeed, her case is one of rare serendipitous malfunctioning with respect to belief-regulation. Here, then, the particular species of epistemic criticism in question is, in the first instance, criticism leveled at the pattern of belief-regulation that underwrites Bella's belief that 2343 is prime, and then (derivatively) at the belief itself.

So, within the action-oriented epistemological framework, there are at least two ways in which an apparently false-but-useful belief might fall short, epistemically speaking. Since each way corresponds to a distinct dimension of epistemic evaluation, so a false belief that facilitates successful action might engender either or both species of negative epistemic evaluation. Thus, the action-oriented epistemologist who accepts that certain false beliefs may be epistemically flawless need not and should not concede that every false belief which facilitates or would facilitate a particular subject's successful action in the particular circumstances she faces is thereby epistemically flawless and so immune from epistemic criticism.

Both action-oriented epistemology and straightforward pragmatism about cognitive evaluation embrace the thought that structure and character of epistemic evaluation is somehow derivative of and grounded in the contribution that beliefs paradigmatically make to the production of successful action. But adopting an action-oriented epistemology does not involve embracing a straightforward or full-scale reduction of epistemic norms to straightforwardly practical, prudential, or pragmatic norms. Beliefs which happen to facilitate successful action in a particular set of circumstances and in the service of a particular set of ends may not be well-suited to facilitate action across the variety of potential circumstances that human subjects normally face and in the service of the variety of different ends they normally have. Moreover, belief-regulating cognitive processes which happen, on one occasion, to generate a belief that is well-suited to facilitate successful action across a variety of different circumstances may not do so routinely, predictably, or reliably. And this means that on the action-oriented approach, epistemic evaluation certainly can, and sometimes will, come apart from practical or prudential evaluation.

³³ It is worth noting that an action-oriented epistemology faces its own version of the generality problem here: how can/should the action-oriented epistemologist identify or pick out the belief-regulating processes that are to be evaluated for generating and sustaining, e.g., Bella's belief? I simply do not have space here to offer a full solution to the action-oriented epistemologist's version of the generality problem. However, I will suggest that empirical work on belief fixation from cognitive science and psychology supplies the action-oriented epistemologist with useful resources for developing a solution to the problem. Indeed, the discussion in the preceding sections of the paper foreshadows this point. If empirical work suggests that the human cognitive system is comprised of relatively encapsulated, domain-specific belief-regulating mechanisms/processes (e.g. belief-regulating mechanisms that come on-line, so to speak, in the domain of mathematics, or in the domain of self-regarding beliefs, and perhaps our beliefs about our friends and loved ones), then the action-oriented epistemologist ought to take this preliminary evidence that these relatively encapsulated, domain-specific mechanisms/processes are appropriate targets for epistemic evaluation.

Is it possible to embrace a more nuanced account of evaluation of Bella's belief without giving up the spirit of straightforward pragmatism? Well, believing in the absence of evidence is *often* practically disastrous. And so, the way in which Bella ignores or disregards her (lack of) evidence in regulating her belief here is risky. Accordingly, then, perhaps the straightforward pragmatist can say that Bella's belief is criticizable by virtue of subjecting Bella to an objectionable (or, perhaps, simply regrettable) level of risk. So, Bella's belief is praiseworthy by virtue of being well-suited to facilitate her actions in the particular (idiosyncratic) circumstances that she faces. And the belief-regulating processes that generate and sustain her belief are praiseworthy by virtue of producing such a belief. But it is notably *lucky* that Bella's belief and the belief-regulating processes that generate and sustain her belief are praiseworthy in these ways—indeed, both results were unlikely, given the general facts about the circumstances in which human believers find themselves. And so perhaps the now-somewhat-less-straightforward pragmatist can say that Bella's belief (and/or the belief-regulating processes that generate and sustain her belief) are criticizable for this reason.

Notice, however, that the now-somewhat-less-straightforward pragmatist's strategy here involves distinguishing at least two different species of what we might think of as broadly practical (or action-oriented) evaluation: one that diagnoses Bella's belief as praiseworthy or virtuous, and another that diagnoses Bella's belief as defective or flawed. But once we embrace this distinction, the distance between somewhat-less-straightforward pragmatism and an action-oriented epistemology shrinks considerably, perhaps to the point of being merely terminological. The real disagreement between the action-oriented epistemologist and the somewhat-less-straightforward pragmatist seems now to be a disagreement about whether a certain species of broadly practiced evaluation is properly classified as *epistemic*. But the action-oriented epistemologist circumscribes the epistemic domain in an intuitively attractive and principled way (recall that the action-oriented approach embraces MC, which states that the structure and character of epistemic evaluation is determined by what beliefs *are*, and so identifies epistemic evaluation as evaluation with respect to those standards or norms that are grounded in and derive from belief's nature). So, without offering an argument against MC, it seems that the somewhat-less-straightforward pragmatist is in no position to reject the action-oriented epistemologist's claim that the relevant sort of criticism is distinctively *epistemic* in character.

The action-oriented approach vindicates the intuitively attractive idea that the epistemic perspective is, in some sense, autonomous. Although the action-oriented approach understands epistemic norms as action-oriented, epistemic evaluation has a distinctive character, and it retains its independence from other, familiar species of straightforwardly practical evaluation. Thus, an action-oriented epistemology is substantially less revisionary than straightforward pragmatism. It involves a less radical departure from the ordinary ways of thinking about the independence of the epistemic evaluation and the possibility of conflict between epistemic evaluation and practical or prudential evaluation that lurk behind our epistemic practice. I take it that, *ceteris paribus*, we ought to prefer approaches to epistemological theorizing that can vindicate the pre-theoretical commitments that underwrite our epistemic practice. Thus, in the absence of some reason to think that a more revisionary approach is required

(and I see no such reason), we ought to prefer the action-oriented approach to its straightforwardly pragmatic competitor.³⁴

References

- Berker, S. (2013). The rejection of epistemic consequentialism. *Philosophical Issues*, 23(1), 363–387.
- Bortolotti, L., & Antrobus, M. (2015). Costs and benefits of realism and optimism. *Current Opinion in Psychiatry*, 28(2), 194–198.
- Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Boston: MIT Press.
- Engel, A., et al. (2013). Where's the Action? The pragmatic turn in cognitive science. *Trends in Cognitive Science*, 17(5), 202–209.
- Engel, A., Friston, K., & Kragic, D. (eds.), *The pragmatic turn: Toward action-oriented views in cognitive science* (pp. 19–34). Cambridge: MIT Press.
- Gibbard, A. (2005). Truth and correct belief. *Philosophical Issues*, 15, 338–350.
- Gibbons, J. (2013). *The norm of belief*. Oxford: Oxford University Press.
- Hazlett, A. (2013). *A luxury of the understanding: On the value of true belief*. Oxford: Oxford University Press.
- Johnson, D. D., & Fowler, J. H. (2011). The evolution of overconfidence. *Nature*, 477(7364), 317–320.
- Keller, S. (2004). Friendship and belief. *Philosophical Papers*, 33, 329–351.
- Kornblith, H. (1993). Epistemic normativity. *Synthese*, 94(3), 357–376.
- Kornblith, H. (2002). *Knowledge and its place in nature*. Oxford: Oxford University Press.
- Lycan, W. G. (1988). *Judgement and Justification*. Cambridge: Cambridge University Press.
- Marušić, B. (2015). *Evidence and agency norms of belief for promising and resolving*. Oxford: Oxford University Press.
- McCormick, M. S. (2014). *Believing against the evidence: Agency and the ethics of belief*. New York: Routledge.
- McKay, R., & Dennett, D. (2009). The evolution of misbelief. *Behavioral and Brain Sciences*, 32, 493–561.
- Millikan, R. G. (1993). Naturalist Reflections on Knowledge. *White queen psychology and other essays for Alice* (pp. 241–264). Cambridge: MIT Press.
- Nolfi, K. (2015). How to Be a Normativist About The Nature Of Belief. *Pacific Philosophical Quarterly*. <https://doi.org/10.1111/papq.12071>.
- Nolfi, K. (2018a). Another kind of pragmatic encroachment. In B. Kim & M. McGrath (Eds.), *Pragmatic Encroachment* (pp. 35–54). Abingdon: Routledge.
- Nolfi, K. (2018b). Why only evidential considerations can justify belief. In C. McHugh, J. Way, & D. Whiting (Eds.), *Normativity: Epistemic and Practical* (pp. 179–199). Oxford: Oxford University Press.
- Nolfi, K. (2019). 'Epistemic Norms. *All Things Considered*'. *Synthese*. <https://doi.org/10.1007/s11229-019-02485-2>.
- Papineau, D. (1987). *Reality and representation*. New York: Blackwell.
- Pezzulo, G. (2015). The contribution of pragmatic skills to cognition and its development: common perspectives and disagreements. In A. Engel, K. Friston, & D. Kragic (Eds.), *The pragmatic turn: Toward action-oriented views in cognitive science* (pp. 19–34). Cambridge: MIT Press.
- Preston-Roedder, Ryan. (2018). Three varieties of faith. *Philosophical Topics*, 46(1), 173–199.
- Quine, W. V. (1969). Natural kinds. In J. Kim & E. Sosa (Eds.), *Ontological relativity and other essays* (pp. 114–138). New York: Columbia University Press.
- Ramsey, F. P. (1927). Facts and propositions. *Proceedings of the Aristotelian Society (Supplementary)*, 7, 153–170.

³⁴ I am grateful to my colleagues in the University of Vermont Philosophy Department, and to audiences at The University of Connecticut, UQÀM, the 2016 Conference on False But Useful Beliefs, the 2016 Eastern Division Meeting of the American Philosophical Association, and the 2015 Penn Reasons and Foundations in Epistemology Conference for their invaluable feedback on this material. I am particularly indebted to Roger White for his comments on an earlier version of this manuscript, and to Ram Neta and Susanna Rinard for many conversations that helped me develop the ideas that find voice here.

- Rinard, S. (2015). No exception for belief. *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12229>.
- Shah, N. (2003). How truth governs belief. *Philosophical Review*, *112*, 447–482.
- Shah, N., & Velleman, J. D. (2005). Doxastic deliberation. *Philosophical Review*, *114*(4), 497–534.
- Shapiro, L. A. (2010). *Embodied cognition*. New York: Routledge.
- Sharot, T. (2011). *The optimism bias: A tour of the irrationally positive brain*. New York: Vintage.
- Sosa, E. (2009). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. I). Oxford: Oxford University Press.
- Stich, S. (1993). *The fragmentation of reason: A preface to a pragmatic theory of cognitive evaluation*. Boston: MIT Press.
- Stroud, S. (2006). Epistemic Partiality in Friendship. *Ethics*, *116*(3), 498–524.
- Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin*, *103*(2), 193–210.
- Taylor, S. E., & Brown, J. D. (1994). Positive illusions and well-being revisited: separating fact from fiction. *Psychological Bulletin*, *116*(1), 21–27.
- Velleman, J. D. (2000). *The possibility of practical reason*. Oxford: Oxford University Press.
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, *16*, 267–297.
- Wedgwood, R. (2007). *The nature of normativity*. Oxford: Oxford University Press.
- Wedgwood, R. (2013). The Right Thing to Believe. In Timothy Chan (Ed.), *The Aim of Belief*. Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, *9*(4), 625–636.

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