

# *Evidentialism in action*

**A. K. Flowerree**


**Philosophical Studies**

An International Journal for Philosophy  
in the Analytic Tradition

ISSN 0031-8116

Philos Stud

DOI 10.1007/s11098-019-01376-z



**Your article is protected by copyright and all rights are held exclusively by Springer Nature B.V.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at [link.springer.com](http://link.springer.com)".**



## Evidentialism in action

A. K. Flowerree<sup>1</sup> 

© Springer Nature B.V. 2019

**Abstract** Sometimes it is practically beneficial to believe what is epistemically unwarranted. Philosophers have taken these cases to raise the question are there practical reasons for belief? Evidentialists argue that there cannot be any such reasons. Putative practical reasons for belief are not reasons for belief, but (to use a distinction from Pamela Hieronymi) reasons to manage our beliefs in a particular way. Pragmatists are not convinced. They accept that some (or perhaps all) reasons for belief are practical. The debate, it is widely thought, is at an impasse. But this debate fails to address what is puzzling and interesting about the cases. By focusing on reasons for belief, the debate completely overlooks the role of action in relation to belief. We should be talking about the reasons for actions that shape our beliefs, which I will call belief management. I argue for three related theses: (1) the interesting cases that motivate the debate are about belief management; (2) Evidentialism is irrelevant to belief management; (3) agents have practical reasons to manage their beliefs with the aim of forming true beliefs. These reasons are categorical in nature and result in the tension of conflict cases.

**Keywords** Ethics of belief · Evidentialism · Pragmatism · Practical reasons for belief · Philosophy of action

---

✉ A. K. Flowerree  
Amy.flowerree@ttu.edu

<sup>1</sup> Philosophy Department, Texas Tech University, Box 43092, Lubbock, TX 79409, USA

## 1 Introduction

### BOOK

Suppose you learn that if you believe you are a brilliant philosopher, you are far more likely to finish writing your first book. Suppose, further, that philosophy is the sort of discipline in which it is very difficult to assess whether someone who has not yet finished her book is a brilliant philosopher. Your evidence for your brilliance is not very clear. But your tenure clock is ticking, and you must finish in order to remain gainfully employed. Should you believe you are brilliant, though your soul is weary and your confidence wearing thin?

BOOK raises a perennial philosophical question. When our evidence suggests we should believe one thing, and our practical interests suggest another, what should we believe? Philosophers have taken cases like BOOK to be asking *are there practical reasons for belief?* Evidentialists argue that there cannot be any such reasons. According to many Evidentialists, I cannot believe for practical reasons because reasons for belief are picked out by their functional role: reasons for belief that  $p$  are considerations that bear on the truth of  $p$ .<sup>1</sup> There cannot be practical reasons for belief because practical reasons cannot play the right role (in showing  $p$  to be true).<sup>2</sup> Putative practical reasons for belief are not reasons *for* belief, but (to use language from Pamela Hieronymi) reasons to *manage* our beliefs in a particular way. Pragmatists are not convinced. They accept that some (or perhaps all) reasons for belief are practical. That is, according to the Pragmatist, there is a nontrivial way in which practical reasons could justify belief (directly, or indirectly). Evidentialists double down with conceptual arguments about how beliefs are justified.<sup>3</sup> Pragmatists reject these arguments as ad hoc.<sup>4</sup> The result is an impasse.

But, I will argue, this debate fails to address what is puzzling and interesting about cases like BOOK. By focusing on technical notions of reasons for belief, the debate completely overlooks the role of action in relation to belief. We should be talking about the reasons for actions that shape our beliefs, which I will call *belief management*. I argue for three related theses: (1) the interesting cases that motivate the debate are straightforwardly about belief management; (2) Evidentialism is irrelevant to belief management; and (3) agents have practical reasons to manage their beliefs with the aim of forming true beliefs. These reasons are categorical in nature and result in the tension seen in cases like BOOK. Thus, the nature of action

<sup>1</sup> Evidentialists give various defenses of this claim. Nishi Shah, for example, argues something could be a reason to  $\phi$  just in case one could  $\phi$  for that reason, (Shah 2006). Since we cannot bring about belief on the basis of practical reasons, practical reasons cannot be reasons for belief. See also Shah (2011). Hieronymi (2008) argues that reasons for belief are reasons that would settle for me *whether p?* Berker (2019) has endorsed a similar picture. See also Wedgwood (2002, 2013) and Scanlon (2014). This is not the strategy of the original evidentialist, Clifford (1877), and I set him aside for this paper.

<sup>2</sup> Shah (2006) and Hieronymi (2005). For a Pragmatist criticism, see Rinard (2015).

<sup>3</sup> Shah (2011).

<sup>4</sup> Rinard (2015).

itself grounds a kind of evidentialism. Evidentialism, I conclude, percolates from action.

## 2 The cases

Cases like Book are of perennial philosophical interest. Pascal's wager is taken as the starting point for discussions. But Pascal doesn't actually present his wager a reason to *believe* in God. Rather, it is presented as a reason to *get oneself by whatever means* to believe in God. Consider this version:

### WAGER

Suppose you read Pascal, and you become convinced of the wager. If God exists and you believe that God exists, you will receive infinite reward; if you fail to believe, you will receive infinite punishment. If God does not exist and you believe, you lose very little; if you fail to believe, you have no gain except avoiding a false belief. You think the evidence for God's existence is inconclusive. Pascal urges you to immerse yourself in a religious community in order to induce the belief in God. Should you do it?

This is straightforwardly about what *actions* you should undertake. Should you organize your life in such a way that you will bring it about that you believe in God? But if Pascal's wager is understood as a justification for whether to *act* a certain way, then the Evidentialist position about reasons for belief does not directly (or indirectly, I will argue in the next section) bear on what to do in cases like WAGER.

People have a lot to say about Pascal's Wager.<sup>5</sup> My point here doesn't ride on the particulars of the wager. So let's consider three other popular cases:

### SELF-ENHANCEMENT

Psychologists reiterate that if you believe yourself to have positive attributes, you are more likely to be successful. Furthermore, if you construct a mantra (e.g. *I am good-looking, capable, and people like me!*) and say it faithfully every time you feel down on yourself, you may be able to form the belief that you have these positive attributes, which will help you be successful. Should you use a mantra?<sup>6</sup>

### DOGMATISM

You are antecedently disposed to oppose a particular government policy because a politician you despise endorses it. You have a friend who expresses enthusiasm about the policy on facebook, and she posts a link about the positive impact of the policy (you think; you don't click on the link). Your impulse is to unfriend her, and read only articles by those who oppose the policy. Is it permissible for you to turn off your access to evidence that might

<sup>5</sup> For an overview of the complexities of Pascal's Wager, see Hájek (2018).

<sup>6</sup> See Hazlett (2013) for a version of this case.

go against positions you already hold, simply because the evidence annoys you?

#### HIGH STAKES

You live on the Iberian Peninsula during the Middle Ages. The conquering Christians offer you the choice: convert or die. You know that in order for your conversion to be convincing, you will need to get yourself to really believe in the Christian God, not merely offer lip service. What should you do, get yourself (by whatever means) to believe in the Christian God, or face death?<sup>7</sup>

There are three things to note about these cases. First, WAGER, SELF-ENHANCEMENT, DOGMATISM, and HIGH STAKES are straightforwardly about actions. We can undertake actions that may result in our having certain belief states. Should we engage in behavior that will induce a belief, regardless of our evidence for the truth of the content of the belief? In order to answer the question raised in these cases, we must understand which considerations bear on how we ought to manage our beliefs.

Secondly, a salient feature of these cases is that they elicit a *tension*. We feel the pull of our practical interests, and we also resist manipulating ourselves into forming these beliefs. While the exact weighing of the reasons may be unclear, there are *pro tanto* reasons on both sides. Even if we are inclined to judge that in SELF-ENHANCEMENT, for example, that it is permissible for the agent to bring about a belief in her superiority, we still sense that there is some reason to pursue an honest self-conception. And even if we are inclined to judge in DOGMATISM that the agent ought to avoid closing off her access to evidence that goes against her beliefs, we can understand the agent who chooses to do so. There is legitimate reason to manage our beliefs to suit our practical interests. There is also reason, I think, for us to manage our beliefs with the aim of gaining truths.

And finally, this tension pulls in uneven ways. In DOGMATISM, The dogmatic person who closes herself off from new evidence because she finds it uncomfortable is not—intuitively—in the same position as the person who closes herself off to protect her emotional wellbeing. It isn't as though practical concerns uniformly trump the concern for true beliefs. An account of reasons for belief management needs to explain this tension.

### 3 The limits of Evidentialism

In this section, I show that Evidentialism is irrelevant to the cases at hand.<sup>8</sup> The Evidentialist begins with an intuitive distinction: there are reasons for belief and reasons for belief management. Reasons for belief, thinks the Evidentialist, are

<sup>7</sup> If you have difficulty seeing this as a real dilemma, at least recognize that many decided to convert, and many decided to die. For a different case of high stakes, see James's Chasm Jumping case (James 1897).

<sup>8</sup> Many views are labeled "Evidentialist." I am concerned with those listed in footnote 1. Conee and Feldman-style evidentialists straightforwardly accept that their version of evidentialism does not extend to action (see Feldman 2000, p. 688–690). Another version of evidentialism might think that practical and epistemic evaluation are simply two different standards of assessment, so that we can evaluate any

considerations that count in favor of the truth of  $p$ . Reasons for belief management are reasons that count in favor of me bringing it about that I believe  $p$ . In BOOK, the fact that a belief in my genius would help me achieve my ends is a reason to bring it about that I believe  $p$ . But it is not a reason to believe  $p$ . And so, the Evidentialist concludes, there are no practical reasons for belief.

Suppose we accept the Evidentialist answer to the question as it is initially formulated: there are no practical reasons for belief. Further, suppose we accept that this is for the reason many Evidentialists favor, because only a consideration that counts in favor of the truth of  $p$  could be a reason for believing  $p$ . It only seemed like there were practical reasons for belief because we failed to distinguish between *reasons for belief* and *reasons for belief management*. While there may be practical reasons to *bring it about* that I believe  $p$ , there are no practical reasons to *believe*  $p$ .

Reasons for belief, on this view, are considerations that the agent takes to be authoritative to settling the question *whether  $p$* ? This question is a first-order question, a question that is primarily world-involving, not agent involving. The considerations brought to bear provide rational support for an *attitude* with a mind-to-world direction of fit. Reasons for belief management, by contrast, embody answers to the question *whether I should bring it about that I believe  $p$* ? This question is a second-order question, a question that involves the agent assessing what state of affairs she should bring about. The considerations that bear on this question rationalize actions, and (plausibly) have a world-to-mind direction of fit.

If we grant the Evidentialist her distinction, it follows that there are no practical reasons for belief. But this does not illuminate the cases that we are interested in. The underlying tension, the one that we face as agents, is not resolved by distinguishing the norms that govern rational belief formation from the norms that govern belief management. Think about BOOK. The fact that believing I'm brilliant will help me finish my book is *not* a reason to believe I am brilliant (on the evidentialist's view). And yet, suppose I ask *whether to induce the belief that I am brilliant*? Perhaps I engage in a mantra (as SELF-ENHANCEMENT suggests). What good is my intellectual humility and self-knowledge if I am an unemployed philosopher? But on the other hand, I was deeply moved by the Delphic Oracle's command to "know thyself." The Oracle's command involves undertaking a certain kind of project, a project of seeking out evidence so that I can form an accurate self-understanding. Should I follow the Delphic Oracle, or hum my mantra and power through?

The Evidentialist has given us a framework for settling first order questions about what to believe. But the agent's relationship to her beliefs is far more complex and demanding than this simplistic framework acknowledges. Agents continually face questions about which epistemic practices to engage in, and these practices have a substantial impact on belief formation. Our belief management directly influences the resources with which we can settle first order questions about what to believe.

---

Footnote 8 continued

attitude or action from the epistemic point of view, or the practical point of view (See Alston 1989, pp. 83–84). While I think similar arguments can be made against this version, it is not my target here.

Even if we accept the Evidentialist picture that there are no practical reasons for belief in a *direct* sense, there certainly are practical reasons for belief in an *indirect* sense.<sup>9</sup> Our cases illuminate this very tension. Even if our epistemic reasons are clear in a given situation, we can manage our beliefs to change our epistemic situation. And we clearly have practical reason to do this sometimes. Perhaps I am a paramedic and the patient needs a blood transfusion. At the moment, I don't know the patient's blood type. Suspension of judgment, while epistemically respectable, is not a professional option for me. I have reason to change my epistemic situation so that I can form a belief about the patient's blood type.

In our target cases, the agent is left with a pressing question, how should I manage my beliefs? What considerations should I take to settle this question? Is it my practical interests? Or should I manage my beliefs with the aim of believing only truths? But what can the Evidentialist say about the problem facing the agent? The agent has a problem that is not resolved by telling her that there are no practical reasons for belief, only practical reasons for belief management. Her reply will then be, incredulously, "Well, then, how should I *manage* my beliefs?" The Evidentialist has only pushed back the problem.

And, as should be apparent by now, the Evidentialist cannot account for the tension of the original cases. There is nothing contradictory between answering "Whether *p*?" affirmatively, and "Whether to manage my beliefs so that I believe *p*?" negatively (or vice versa). Philosophers have used cases like Book to ask, "Are there practical reasons for belief?" The Evidentialist resolves this question by showing that there cannot be. But this answer is irrelevant to the agent struggling to decide what to do. For her, the question is, "How should I go on with respect to my beliefs?" The Evidentialist can give no insight into this question.

Now we are in a position to diagnose why Evidentialism is such an unsatisfying position. The Evidentialist evades, rather than answers, the central question of cases like Book. Agents have complex relations to their beliefs, and our central cases are probing exactly what reasons govern these relations. The Evidentialist picture does not have the resources to capture the tension of our cases. And even if the Evidentialist is right about reasons for belief, the agent's predicament about what to do is entirely unresolved.

One might think that the Pragmatist is the default champion, but this would be too quick. Pragmatism, as we set up the debate, is just the negation of Evidentialism. It is the thesis that there are at least some practical reasons for belief. But once we shift the question to reasons for action, it is unclear what the distinctively "pragmatist" view would be. In what follows, I explore the nature of belief management in order to clarify our target question. In the final section, I reconstitute the debate as between the Instrumentalist—who holds that belief management should be governed by what is instrumental to achieving our ends—and the

<sup>9</sup> Rinard (2015) and Leary (2017) both make a similar point. They take this to show that the distinction between reasons for belief and reasons for belief management is superficial. I disagree with this conclusion; I think the distinction is conceptually and normatively important. Instead, I think this point illuminates that the normative pressure on the agent is *fraught*. More on this in the next section.



Aletheist—who holds that belief management should be governed by the aim of producing true beliefs. I will then give an argument in favor of Aletheism.

#### 4 Belief management

Our actions have an enormous impact on what we end up believing. Arguably, *every* action has some impact on what we end up believing. But for now, I am interested in a subset, those actions undertaken with the aim of shaping our beliefs. Let's call this set of actions *belief management actions*. Belief management is introduced by Pamela Hieronymi in "Believing at Will."<sup>10</sup> In this paper, she distinguishes between an epistemic question (*whether to believe p?*), and a practical question, (*whether to bring it about that I believe p?*). When we aim to bring about, or induce within ourselves, a belief that *p*, we are engaging in belief management. Hieronymi leaves the distinction at an intuitive level.<sup>11</sup> Hieronymi suggests that the question of whether to induce the belief that *p* is parallel to the question of whether to make soup.<sup>12</sup> We can make soup for any reason we take to support soup-making. Similarly, we can induce the belief that *p* for any reason we take to support inducing the belief that *p*. While Hieronymi doesn't elaborate, the idea is that the reasons to induce a belief are the same sort as the reasons that determine whether to make soup: i.e., is it a good thing to bring about? Once we settle the answer to this question, we can then formulate a plan of action to bring about the desired result.

Belief management involves my stance towards my beliefs. It involves, on the one hand, actions I might take that bring me into contact with evidence (or help me avoid evidence). And, on the other hand, it involves judgments about deliberation. Belief management involves determining when to open deliberation, how to direct my attention while I deliberate, and when to terminate deliberation.

We engage in belief management all the time. In the standard case, we manage our beliefs by gathering evidence and initiating first order deliberation. If I am thinking about what to make for dinner, I might decide to go look in the fridge to see what is available. If I am an employer looking to hire, I might realize that when I'm tired I tend to form biased beliefs against minority job candidates. As a result, I might structure my inquiry in a way to improve my chance of success, by looking at minority applicant files more closely, or at the beginning of a review session. These examples involve working with rational belief forming processes. But my belief management could aim at truth, even if I recognize the first order belief forming process has been compromised. I might have a recalcitrant belief that my hands are dirty (despite the fact that I have washed them again and again) and engage in a

<sup>10</sup> Hieronymi (2011). The distinction is implicit in Hieronymi (2005).

<sup>11</sup> She is primarily interested in the kind of agency we display when forming beliefs and intentions, in contrast to the kind of agency we display when we act. For Hieronymi, the distinction illuminates two kinds of control, one she calls evaluative control (which is the kind of control we enjoy over our beliefs) and managerial or manipulative control (the kind we enjoy over ordinary objects). See Hieronymi (2008). Here, Hieronymi calls it manipulative control. See also Hieronymi (2006).

<sup>12</sup> Hieronymi (2011, 17).

mantra to help me remind myself that I am not overwhelmed with germs. These examples illuminate belief management that is truth-sensitive. We engage in actions that are aimed at inducing true beliefs. We gather evidence, structure our inquiry, and even sometimes bypass rational deliberation in pursuit of gaining true beliefs. Let's call this kind of belief management—belief management that aims at forming true beliefs—Alethic Belief Management.

Practical considerations are always present in our belief management. Importance and time-sensitiveness determine how hard we work to make sure our beliefs come out true. But even though importance and time-sensitiveness limit our effort, our belief management is generally not structured by practical considerations. Our aim is generally to end up with true beliefs. So while practical considerations constrain the lengths to which we will go to end up with true beliefs, the aim of belief management is standardly the formation of true beliefs. Descriptively, our practice is consistent with two general aims we might have.

*Alethic belief management* actions that aim at producing epistemically good beliefs.<sup>13</sup> Paradigmatic instances of Alethic Belief Management include seeking out evidence, opening deliberation, structuring deliberation to minimize bias, judging when to close deliberation, and even overriding deliberation known to be biased.

*Instrumental belief management* actions that aim at producing beliefs for some other end, without regard for their epistemic status. Paradigmatic instances of Instrumental Belief Management include circumventing first order deliberation in order to maintain a biased belief, exploiting cognitive biases in order to form or maintain a belief regardless of its truth, and causally/psychologically manipulating ourselves. If there were any, it could include belief-inducing pills.

Alethic Belief Management unifies the aim of belief management with the first order question of *whether p?* While practical issues such as timeliness and importance will always be present in our belief management, Alethic Belief Management illuminates a fundamental structural connection between the actions we take that foreseeably bring about belief, and the epistemic standing of the beliefs themselves. Our practical interests will dictate where we direct our attention, when we decide to deliberate over *whether p?*, and how much time we have to collect evidence. Alethic Belief Management unifies the agent's view towards her own beliefs; the agent views both her beliefs and her belief management to share an aim, namely the aim of bringing about true beliefs.

Exploiting cognitive biases or causally manipulating ourselves need not be non-alethic. If I come to realize that my first order deliberative capacity is compromised, then I might take steps to correct for this deficiency. For example, I might fail to appreciate someone's testimony because they belong to a disadvantaged group, resulting in what Miranda Fricker calls a credibility deficit.<sup>14</sup> My failure to appreciate another's testimony arises from my first order capacity being compromised. But perhaps there's nothing directly I can do in the moment to fix this

---

<sup>13</sup> I use "alethic" as a neutral way of characterizing a close connection to truth. The precise status may be truth, knowledge, justification, understanding, or some further thing.

<sup>14</sup> See Fricker (2007).

capacity. In the moment, the best I can do is take steps to circumvent this capacity in order to bring my judgment in line with the truth. And this can only be achieved by managing my beliefs.

By contrast, Instrumental Belief Management is not continuous with first order reasons for belief. Instrumental Belief Management requires that we hide the first-order considerations that usually guide our belief formation. Such activity requires either self-deception or physical/psychological manipulation. We cannot clearly form beliefs at will.<sup>15</sup> For example, we must exploit cognitive biases or take belief-inducing pills (if they can be found). But Instrumental Belief management is unified with our general orientation of setting ends and pursuing them. If our ends require giving ourselves certain beliefs, then an instrumentalist principle would say that we are required to give ourselves certain beliefs (regardless of their truth).

Alethic and Instrumental Belief Management are distinguished by their *aim*. Alethic Belief Management aims at producing true beliefs; Instrumentalist Belief Management aims at producing beliefs for some end that is insensitive the truth, such as practical benefit. An example of this view of Belief Management would be to aim to manage my beliefs so that I have the beliefs that maximizes expected utility.<sup>16</sup> It could also be an aim to manage beliefs to protect a fragile ego, please an authority figure, satisfy a perceived moral requirement, or save my life. On this second view, reasons for belief management are conditioned on our ends, and there force is instrumental.

Here I want to distinguish between an agent that sets an aim (e.g. maximize expected utility), where the pursuit of this aim directs her to manage her beliefs for truth and an agent that makes it her aim to directly manage her beliefs according to a non-alethic aim (e.g. what will maximize expected utility). It could be that these two agents often manage their beliefs the exact same way. The key distinction between these two cases is that in the first case the agent uses alethic belief management to pursue her aim. In the second case, the agent aims to manipulate her beliefs according to her aim. These are two distinct aims, aims that will, depending on circumstances, generate different actions and result in different beliefs.<sup>17</sup> The cases that I am interested in in this paper are the cases where those two aims come into conflict.

From what I've said so far, it is open whether agents should engage in alethic or instrumental belief management. While alethic belief management is unified with first order believing, we do not yet have any justification for why the agent should pursue this aim. And while non-alethic belief management requires self-deception or manipulation, we do not yet have any justification for why that would be unjustified for the agent to do, given the instrumental benefits.

With this framework in hand, let's consider what is happening in our cases. WAGER, SELF-ENHANCEMENT, DOGMATISM and HIGH STAKES ask whether it is rationally

<sup>15</sup> Bennett (1990), Williams (1973) and Hieronymi (2005, 2009a, b).

<sup>16</sup> An example of this view would be (Rinard 2015).

<sup>17</sup> This is true even if, in some cases, the aims recommend the same actions and the same beliefs. While I can't treat the issue fully here, the distinction is roughly the one that (White 2010) examines between features that are relevant for the epistemic status of belief, and those that are irrelevant.

permissible in certain circumstances to engage in Instrumental Belief Management. In *WAGER*, you consider whether you should follow a course of action that will induce a belief in God's existence, and your reasons for adopting this course of action is that it will maximize your expected utility. In *SELF-ENHANCEMENT*, you consider whether to engage in a mantra to induce a belief in your abilities, so that you will be more successful in undertaking difficult projects. In *DOGMATISM*, you are tempted to close off your access to new information because you do not want to challenge your existing beliefs. In the example, you are motivated to do this because you find it unpleasant to be confronted with counter-evidence. In *HIGH STAKES*, you are pulled to take whatever action you can to make sincere your religious conversion, so that you can save your life. You have, it seems, practical reason to engage in Instrumental Belief Management.

But the agent also feels the pull of managing her belief for truth. So, I suggest, the tension in our original cases arises from the agent feeling pulled in two directions. On the one hand, she is pulled to engage in Instrumental Belief Management, for the benefits she can bring about. And on the other hand, she is pulled to engage in Alethic Belief Management.

In these cases, we can clearly see the reasons for Instrumental Belief Management. But what reason would there be—in these cases—for Alethic Belief Management? In the final section, I will argue that all agents have categorical reason to pursue Alethic Belief Management. In my view, the tension is between an agent bringing about her ends, and the agent acting rationally. My argument is not intended to exhaust possible explanations of the tension. Rather, my argument shows that the concern for truth arises from the nature of practical agency itself.

## 5 Evidentialism in action

So far, I have argued that the target cases are really about reasons for belief management, not reasons for belief. I have also shown that the Evidentialist position doesn't have anything to say about reasons for belief management. I then characterized belief management and sorted it into two types, distinguished by their aim, Alethic and Instrumental. In this final section of the paper, I reconstitute the debate in its proper domain, as a debate over reasons for belief management. I will now argue that agents have *pro tanto* reason to engage in alethic belief management, regardless of the context, aims, or stakes. Let's call this view Aletheism.

**Aletheism** Practically Speaking, the beliefs an agent ought to have are those that are epistemically well-founded and conducive to achieving her ends  
The foil to my view is Instrumentalism:

**Instrumentalism** Practically speaking, the beliefs an agent ought to have are those that are conducive to achieving her ends

While the Instrumentalist will—of necessity!—reject my conclusion, my aim is to give a presumptive argument in favor of Alethic Belief Management. I will then sketch where I see the debate going from here.

From Instrumentalism it follows that the agent only has reason to engage in Instrumentalist Belief Management. Any reasons we have for belief management will depend on the aims, context, and stakes for the agent. There is no reason to engage in Alethic Belief Management. Instrumentalist Belief Management will often look the same in practice as its Alethic counterpart, since the truth is often beneficial, but when our interests are in tension with managing our beliefs for truth, the agent can have no reason to choose Alethic Belief Management instead. Various philosophers have endorsed something like the Instrumentalist position.<sup>18</sup>

The Instrumentalist thinks that the epistemic grounds of our beliefs matter only when they help the agent achieve her ends. On this view, any reason we have to manage our beliefs for truth is purely instrumental, not alethic. But, I will show, the epistemic grounds of our beliefs matter a great deal, even when they don't help us achieve our ends. Practical reasoning itself requires alethic belief management. The Instrumentalist can avoid defeat only by claiming that practical reasoning also matters only when it helps us achieve our ends. This is tantamount to claiming that practical reasoning places no normative requirements on us. This, I will argue, is counter to our practices of criticizing, holding responsible, and providing rationalizing explanations. Thus, we should either accept Aletheism, or we must settle for a radical (seemingly benighted) view of what agents have reason to do. Or so I will now argue.

My central claim, which I will now defend, is that the nature of practical reasoning gives us categorical reason to engage in Alethic Belief Management. By “categorical” I mean that the reason does not depend on context, stakes, or particular aims. These reasons are *pro tanto*. They may be outweighed, but they are not defeated by the agent's aims or the action's outcomes. In contrast, the Instrumentalist thinks that reasons to manage belief for truth are at most *prima facie*, defeasible by context, stakes, or particular aims. My argumentative strategy is to argue, from plausible, widely held premises, to the conclusion that any agent has reason to engage in Alethic Belief Management. Since this argument holds regardless of the agent's context, stakes, or particular aims, I conclude that any agent has such a reason.

My argument relies on a few assumptions about the nature of belief and the nature of practical reason. The first assumption I make is about the connection between belief and action. Belief plays a central role in practical reasoning. It is not necessary for my purposes that belief is the *only* attitude that plays a role in practical reasoning. Instead, I endorse a weaker principle:

**Belief in Action:** If S comes to believe *p*, then thereby *p* is added to her set of available premises for practical reasoning.<sup>19</sup>

<sup>18</sup> Hazlett (2013) is an example, as well as Rinard (2015, 2017)

<sup>19</sup> For a defense of similar principles, see Schroeder (2010), Locke (2015) and Whiting (2014).

Belief in Action is a descriptive claim about the relationship between belief and action.<sup>20</sup> It is not a claim about how *S* *should* reason. It is a claim about what is available to *S* when she engages in practical reasoning. What is available to her, Belief in Action says, includes the set of propositions she believes. Premises for practical reasoning do not magically pop into existence the moment the agent begins deliberation. They embody her collective outlook on the world. Otherwise, how could she call them into being at the moment of deliberation? Belief in Action also leaves open that other entities could play a role in the agent's practical reasoning. For example, her desires or ends could also serve as premises in her practical reasoning.

Instrumentalists should like Belief in Action. This is because, in many of the cases, the rationale for why you should (practically speaking) believe *p* is because of the role that *p* can then play in your practical reasoning. For example, in SELF-ENHANCEMENT, part of the reason you ostensibly ought (practically speaking) to get yourself to believe in your magnificence is because this belief will enable you to do some practical reasoning. For example, "Should I apply for this prestigious fellowship? Of course!"

The second assumption is this. Just as belief is constitutive of action, so *justified* belief is constitutive of *justified* action. I will limit my focus to the case of practical reasoning. In order for practical reasoning to result in justified action, the agent must have good epistemic grounds for her premises. That is, there is an epistemic norm to practical reasoning. I will call this activity—the activity of relying on premises within one's practical reasoning—*premising*. There are two sides to this. First, in order for the reasoning to be genuinely good reasoning, the agent must have good epistemic grounds. Call this the objective grounds. But secondly, in order for the reasoning to *even count at all* as practical reasoning, the agent must be in a position to take herself to have good epistemic grounds. Call this the subjective grounds.

This second assumption is part of an attractive view of the nature of normative appraisal, but it is not uncontroversial. I cannot exhaustively defend this thesis here, but I will offer some considerations in its favor. When we are reasoning about what to do, we cannot use anything we please as a premise in that reasoning. For example, sometimes an agent uses premises that are epistemically unjustified in his practical reasoning. Suppose that Barnard is deciding whether to take a trip to Paris. He reasons that if he goes to Paris, he will find love and happiness because Paris is the city of love. Barnard believes this of Paris because he has been reading romance novels of questionable accuracy. Barnard decides to go to Paris for this reason. Barnard's reasoning is criticizable, and it is criticizable because the premise he uses to decide what to do is epistemically unjustified. There is a distinctively epistemic flaw to his practical reasoning. In the case of Barnard, his belief lacks what I called objective grounds.

Furthermore, an agent cannot reason practically (either from his perspective, or from ours) if she does not meet a minimal epistemic requirement. Suppose, like

---

<sup>20</sup> It is a commonplace that belief is a constituent of action. See for example, Davidson's seminal piece where he suggests that an action is caused by a belief/desire pair (Davidson 1963).

Barnard, Cai is deliberating about whether to go to Paris. But Cai merely *wishes* that Paris is the city of love.<sup>21</sup> Cai wants to find love and happiness, and he wishes that he would find it in Paris. But he does not *believe* it in any sense. He does not *hope* that it is the city of love.<sup>22</sup> Rather, he just wishes it. Suppose Cai, on the basis of this, goes to Paris. Cai is acting on a whim. He is following his fancy. But he is not acting on the basis of practical reasoning. What he is doing should not appear to us, or to himself, as practical reasoning. Cai's reasoning fails on subjective grounds. Unlike Bernard (who lacked objective grounds and was thus criticizable), Cai fails to engage in practical reasoning at all. In order for it to be practical reasoning, the agent must at least take the premises to be governed by some epistemic constraints.

These considerations suggest that good practical reasoning requires that the premises of our practical reasoning meet certain epistemic standards. Bernard and Cai illustrate the way in which the epistemic considerations enter into both first person and third person evaluations of practical reasoning. Good practical reasoning requires not only that the agent's premises in practical reason are well-founded, but also that the agent take her premises to be well-founded (or be in a position to take them to be well-founded). Practical reasoning does not take place in isolation from or in opposition to epistemic reasoning. Rather, the two are closely knit together. We should accept the following requirement:

**Epistemic Requirement on Practical Reasoning:** S ought to make it the case that S premises  $p$  only if S has good epistemic grounds for  $p$ .

The term "good epistemic grounds" is neutral between the many ways that philosophers articulate the epistemic requirement on practical reasoning. Philosophers have defended various versions of this principle.<sup>23</sup> While there is much debate

<sup>21</sup> Aristotle noted that wishing is not subject to epistemic constraints in the way that practical reasons are. In *Nicomachean Ethics*, III.ii, Aristotle reflects that one can wish for the impossible, but one cannot decide to *do* something that one believes is impossible (Aristotle and Irwin 1999).

<sup>22</sup> Hope may be viewed as a doxastic attitude. It is the shadow between belief and disbelief, and it can be rational or irrational, proportional to the likelihood of  $p$ . Hopes grow dim as the likely hood of  $p$  diminishes, and are irrational when there is no possibility of  $p$ .

<sup>23</sup> Locke (2015) argues that one may premise  $p$  iff one is practically certain that  $p$ . Jason Hawthorne and Stanley (2008) argue that you can premise  $p$  iff you know  $p$ . See also Williamson (2000), Stanley (2005), and Smithies (2012). Fantl and McGrath (2009) only endorse half of the biconditional, arguing that if you know  $p$ , you may premise  $p$ . Lackey (2010) argues that one should premise  $p$  just in case it is reasonable to believe  $p$ . Whiting (2014) argues that the aim of belief is to believe  $p$  only if  $p$  is a practical reason. Whiting's view has the upshot of uniting practical reasons and epistemic reasons, though he approaches it from the standpoint of the aim of belief, rather than as a constitutive feature of practical reasoning. Way and Whiting (2016) defend a similar claim. Bok (1997) argues that when S engages in practical reasoning, S must believe that her premises are true (otherwise her reasoning is unsound, by her own lights).

over *which* epistemic requirement we should accept, it is generally agreed that there *is* such a constraint.<sup>24</sup>

The Epistemic Requirement on Practical Reasoning shows why there is a tension in our original cases. In what follows, I will argue that the epistemic requirement on practical reasoning generates categorical reason to engage in alethic belief management. By practical reasoning, I mean reasoning that aims at determining action (resulting in either a judgment, an intention, or the action itself).<sup>25</sup> All *oughts* in this argument are *oughts* of practical reason.

## 6 Categorical reasons for alethic belief management

(Motivating Assumption)

*Epistemic Requirement on Practical Reasoning (ERPR)* S ought to make it the case that S premises  $p$  only if S has good epistemic grounds for  $p$ .

1. In order to satisfy (ERPR), S has reason to see to it that if  $p$  is added to her set of available practical reasons, then S has good epistemic grounds for  $p$ .
2. *Belief in Action* If S comes to believe  $p$ , then thereby  $p$  is added to her set of available practical reasons (the relationship is entailment)
3. So, in order to satisfy (ERPR), S has reason to see to it that if S comes to believe  $p$ , then S has good epistemic grounds for  $p$ .

ERPR motivates premise (1). First, in the moment of deliberation, the agent cannot be required to investigate whether her available premises have good epistemic grounds. Otherwise, she would set off an infinite chain of deliberations into the epistemic grounds of each of her premises.<sup>26</sup> Action could never take place.

<sup>24</sup> See, for example, Brown (2010). While Brown offers powerful arguments against particular versions of the epistemic requirement, she accepts that there is *some* epistemic requirement, even if it isn't unitary. See also Alvarez (2010). Maria Alvarez argues that the difference between theoretical and practical reasoning is not the form or the premises, but the aim for which it is undertaken. Theoretical reasoning is about what is the case, and practical reasoning is about what to do, but both share the same form and premises. The difference is that practical reasoning results in an action (or the formation of an intention to act). Thus, Alvarez argues, practical reasoning is subject to the same epistemic standards as theoretical reasoning. A notable exception to this remarkable agreement is Parfit (2011, pp. 112–113). See Mueller (2017) for an argument against Parfit's view. See Robertson (2011) for a general argument in defense of an epistemic requirement on practical reasoning.

<sup>25</sup> Here, I wish to remain neutral on the precise output of practical reason. I myself am drawn to Alvarez and Hieronymi's picture, that practical reasoning involves answering a question, "what should I do?" See also Anscombe (1957). But nothing in my argument hinges on this.

<sup>26</sup> This point is made by Aristotle, see *Nichomachean Ethics*, Book III, III3a. Aristotle writes, "For the end cannot be a subject of deliberation, but only the means; nor indeed can the particular facts be a subject of it, as whether this is bread or has been baked as it should; for these are matters of perception. If we are to be always deliberating, we shall have to go on to infinity."

One might object that of course we *can* evaluate in the moment of deliberation. Suppose I deliberate about what to do, and the conclusion I come to is that I ought to punch myself in the face. I might pause here and think, "Surely something went wrong in my deliberation; I better check again." My point is that while it may be possible (as a reasoning failsafe) to pause and rethink, this is a double-checking mechanism, not a first checking mechanism. If it were the first check, Aristotle's point stands. Reflective



Instead, she must *rely* on the epistemic status of her available premises. Thus, the only way she can satisfy (1) is by ensuring from the outset that all her available premises satisfy the epistemic requirement. Second, the purpose—the functional role—of the premises used in practical reasoning are to *justify* actions. That is what the agent is doing when she engages in practical reasoning. So anything that is added into the set of practical reasons ought to be able to fulfill its role. Anything we come to believe, either directly or through belief management, ought to meet the epistemic requirement.

The tight connection between practical reasoning and alethic belief management explains the tension of cases like *Book*. There are reasons to engage in non-alethic belief management. But there are also reasons—reasons generated from practical reason—to engage in alethic belief management. When the agent feels perplexed in this situation, she is responsive to the reasons that are present. There is reason to manage her beliefs for truth. There is reason to manage her beliefs for what is practically beneficial. The tension is generated by two legitimate and opposing normative forces at work within practical reasoning. Believing truths can be in conflict with our practical interests. But it is also required by our practical interests.

If my argument succeeds, then we have a categorical practical reason to engage in alethic belief management. This reason is categorical because the nature of practical reasoning requires our beliefs to meet an epistemic requirement. As a result, it applies to any belief management action, with respect to any proposition, in any context. A concern for the epistemic status of our beliefs originates from practical reasoning itself. We should accept Aletheism.

A reason could be categorical without being overriding. My conclusion is that there is always a reason for alethic belief management. But I have not yet said anything about the relative weight of this reason. Consider the dilemmas of our target cases. How should we weigh the reasons for and against alethic belief management? We have several options. First, the reasons for and against could form a true dilemma. The two kinds of reasons are incommensurable and the agent is doomed to irrationality. Second, it could be that the practical implications override the reasons for alethic belief management. Third, it could be that the reasons for alethic belief management override the instrumental reasons to contravene it. Finally, it could be that the reasons for alethic belief management override some *imes*, but not others.

My aim in this paper is to show that there are categorical reasons for alethic belief management. I am not offering a complete theory of how reasons combine to produce *ultima facie* reasons. But here are just a few considerations for thinking of how to weigh these reasons against each other. Non-alethic belief management undermines your ability to reason and act in a justified way, both by your own lights and by the lights of others. But sometimes it may be *ultima facie* rational to perform this act, over another act that violates an even more fundamental requirement of rationality.

---

Footnote 26 continued

reasoners have multiple ways to intervene on reasoning to double check, but this is not the primary or required way to go. Thanks to Daniel Singer for raising this objection.

Consider Ulysses tying himself to the mast and plugging his ears to avoid the call of the Sirens. Here, Ulysses limits his agency in order to preserve his life and the lives of those aboard his ship. Ordinarily, he has good reason to act in such a way to maintain his physical freedom. But he knows that when he hears the call of the Sirens, he will steer the ship into the rocks and all on board will perish. So in this case, it he has *ultima facie* reason to bind himself in order to survive. Undermining one's agency in a particular situation is sometimes necessary to preserve one's agency overall.

In contrast, consider SELF-ENHANCEMENT. We might think that it is similar to Ulysses. A false self-conception is more effective at promoting one's overall agency, even though it locally limits the agent. But I don't think this stands up to scrutiny. We often describe such people as delusional, narcissistic, and self-aggrandizing. In many cases, an *accurate* self-portrayal will do an even better job, across situations, of helping you meet your goals.<sup>27</sup> It can do this, in part, by helping you set realistic goals. And there is something nearby to self-enhancing beliefs that can play a similar psychological role. A mantra can be rational as the expression of a goal or aspiration, but need not have the goal of inducing a delusional belief. So, it seems to me, SELF-ENHANCEMENT is offering a false dilemma. There is a nearby alternative—setting a regulative ideal, rather than aiming to induce a belief—that will perform the normative work, without the normative cost.

We might wonder how much my argument has moved the dial in the debate over practical reasons for belief. Here, I think, are the notable shifts. First, the pragmatist has gloried in the argument, “if epistemic rationality isn't in your best interests, then so much the worse for epistemic rationality!” This move is no longer available. If my argument is persuasive, the debate has shifted from two different kinds of reasons (epistemic and practical) to a single domain (reasons for action). The debate is not over two conflicting domains of reasons, but the proper account of one domain.

Second, we should abandon the debate over evidentialism and pragmatism (at least as far as thinking about the target cases are concerned). Instead, we ought to be debating Instrumentalism and Aletheism. This is a debate over the nature of practical reasoning, rational action, and agency. If my central argument in this section is cogent, it challenges Instrumentalist views, notably Robust Pragmatism. All reasons cannot be reasons to maximize expected utility. Some reasons are given to us by the structure of reasoning itself. The Robust Pragmatist will likely reject the epistemic requirement on practical reasoning. I think that this will make it difficult to account for our practices of criticizing and blaming those who violate this requirement. But that is a debate for another day.

And finally, even if we accept Aletheism, there are still further questions about how to account for the various cases. This involves exploring (for example) how and when *pro tanto* reasons combine to form *ultima facie* reasons, whether one ought to forgo rationality in order to pursue a greater goal, and whether one's survival is more important than one's rationality.

---

<sup>27</sup> See Young (2014) and Colvin and Block (1994).

## 7 Conclusion

Our target cases (like Book) probe what sorts of reasons agents have to manage their beliefs. Should we manage our beliefs for truth (Aletheism)? Or are beliefs merely instrumental to achieving our ends (Instrumentalism)? I have shown that the Evidentialism vs Pragmatism debate—generally taken to explain our target cases—is irrelevant to these questions. Instead, we ought to focus on the nature of practical reasoning. Practical reasoning, I argued, gives us categorical reason to manage our beliefs so that they end up true.

Once we reframe the debate as about the nature of rational action, fecund lines of inquiry come into view. Is there an epistemic requirement on practical reasoning? How do various *pro tanto* reasons to combine into what the agent ultimately ought to do? How closely related is epistemic rationality and practical rationality?

Finally, the close connection between the epistemic status of our beliefs and rational action reveals the possibility that epistemic reasons have their normative force *in virtue* of the role those reasons play in action. Rather than viewing epistemic reasons as *sui generis*, or reducing them to instrumental reasons, we can articulate a powerful and unified account of reasons that is grounded in a single source, rational agency. On this view, epistemic reasons are constructed from the nature of rational agency, for the purpose of providing inputs to practical reasoning. The resulting view, *epistemic constructivism*, might also shed light on other seemingly intractable debates involving the nature and authority of epistemic norms.

**Acknowledgements** This paper received generous feedback at many venues, and I am grateful to more people than I can name. I am especially thankful to Baron Reed, Selim Berker, James Fritz, Jessica Wright, Luis Rosa, Andy Mueller, Andrea Robitzsch, Lisa Benossi, Jakob Ohlhorst, Cory Davia, and an anonymous referee for formative conversations that shaped this paper.

## References

- Alston, W. (1989). *Epistemic justification*. Ithaca: Cornell University Press.
- Alvarez, M. (2010). *Kinds of reasons: An essay in the philosophy of action*. Oxford: Oxford University Press.
- Anscombe, G. E. M. (1957). *Intention*. Cambridge: Harvard University Press.
- Aristotle, C., & Irwin, T. (1999). *Nicomachean ethic*. Indianapolis: Hackett Publishing.
- Bennett, J. (1990). Why is belief involuntary? *Analysis*, 50(2), 27–107.
- Berker, S. (2019). Combinatorial argument against practical reasons for belief. *Analytic Philosophy*, 59, 427–470.
- Bok, H. (1997). *Freedom and responsibility*. Princeton: Princeton University Press.
- Brown, J. (2010). Knowledge and assertion. *Philosophy and Phenomenological Research*, 81(3), 549–566.
- Clifford, W. K. (1877 [1999]). The ethics of belief. In T. Madigan (Ed.), *The ethics of belief and other essays* (pp. 70–96). Amherst, MA: Prometheus.
- Colvin, R., & Block, J. (1994). Do positive illusions foster mental health? An examination of the Taylor and Brown formulation. *Psychological Bulletin*, 116, 3–20.
- Davidson, D. (1963). Actions, reasons, and causes. *Journal of Philosophy*, 60(23), 685–700.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60(3), 667–695.
- Fricker, M. (2007). *Epistemic injustice*. Oxford: Oxford University Press.

- Hájek, A. (2018). Pascal's Wager. *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), E. N. Zalta (Ed.). Retrieved November 14, 2019 from <https://plato.stanford.edu/archives/sum2018/entries/pascal-wager>.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105, 571–590.
- Hazlett, A. (2013). *A luxury of the understanding: On the value of true belief*. Oxford: Oxford University Press.
- Hieronymi, P. (2005). The wrong kind of reason. *Journal of Philosophy*, 102(9), 437–457.
- Hieronymi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, 87, 45–74.
- Hieronymi, P. (2008). Responsibility for believing. *Synthese*, 161(3), 357–373.
- Hieronymi, P. (2009a). Two kinds of agency. In L. O'Brien & M. Soteriou (Eds.), *Mental action* (pp. 138–162). Oxford: Oxford University Press.
- Hieronymi, P. (2009b). Believing at will. *Canadian Journal of Philosophy*, 39, 149–187.
- Hieronymi, P. (2011). Reasons for action. *Proceedings of the Aristotelian Society*, 111(January), 407–427.
- James, W. (1897). The sentiment of rationality. In W. James (Ed.), *Will to believe and other essays in popular philosophy*. New York: Longmans, Green, and Co.
- Lackey, J. (2010). Acting on knowledge. *Philosophical Perspectives*, 24(1), 361–382.
- Leary, S. (2017). In defense of practical reasons for belief. *Australasian Journal of Philosophy*, 95(3), 878–895.
- Locke, D. (2015). Practical certainty. *Philosophy and Phenomenological Research*, 90, 72–95.
- Mueller, A. (2017) How does epistemic rationality constrain practical rationality? *Analytic Philosophy*, 58(2), 139–155.
- Parfit, D. (2011). *On what matters*. Oxford: Oxford University Press.
- Rinard, S. (2015). Against the new evidentialists. *Philosophical Issues*, 25(1), 208–223.
- Rinard, S. (2017). No exception for belief. *Philosophy and Phenomenological Research*, 93(3), 121–143.
- Robertson, S. (2011). Epistemic constraints on practical normativity. *Synthese*, 181, 81–106.
- Scanlon, T. M. (2014). *Being realistic about reasons*. Cambridge: Harvard University Press.
- Schroeder, M. (2010). *Slaves of the passions*. Oxford: Oxford University Press.
- Shah, N. (2006). A new argument for evidentialism. *Philosophical Quarterly*, 56(225), 481–498.
- Shah, N. (2011). Can reasons for belief be debunked? In A. Reisner & A. Steglich-Petersen (Eds.), *reasons for belief*. Cambridge: Cambridge University Press.
- Smithies, D. (2012). The normative role of knowledge. *Nous*, 46(2), 265–288.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Clarendon Press.
- Way, J., & Whiting, D. (2016). If you Justifiably Believe you ought to  $\Phi$ , then you ought to  $\Phi$ . *Philosophical Studies*, 173(7), 1873–1895.
- Wedgwood, R. (2002). The aim of belief. *Nous*, 36(s16), 267–297.
- Wedgwood, R. (2013). Doxastic correctness. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 87(January), 217–234.
- White, R. (2010). You just believe that because. *Philosophical Perspectives*, 24(1), 572–615.
- Whiting, D. (2014). Reasons for belief, reasons for action, the aim of belief, and the aim of action. In C. Littlejohn & J. Turri (Eds.), *Epistemic norms*. Oxford: Oxford University Press.
- Williams, B. (1973). Deciding to believe. In B. Williams (Ed.), *Problems of the self* (pp. 136–151). Cambridge: Cambridge University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Young, M. C. (2014). Do positive illusions contribute to human well-being? *Philosophical Psychology*, 27(4), 536–552.

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