

# Performance Essentialism

Allan Hazlett

allanhazlett@gmail.com

This paper describes a problem of epistemic normativity, distinct from the ‘value problem’ (§1.1), which I call the problem of the sources of epistemic normativity (§1.2). I lay out a number of possible solutions to the problem (§1.3), including one species of solution that I call performance essentialism (§1.4), based on the idea that epistemic normativity is a species of performance normativity. I go on to criticize three versions of performance essentialism (§2). In my view, the performance essentialist faces an intractable dilemma (§2.1): either she must concede that epistemic evaluation does not apply to believing in general, or else she must say, implausibly, that believing requires truth-directed intentions.

## 1. The sources of epistemic normativity

### 1.1 *The value problem*

Consider what is called the ‘value problem’ in epistemology. The problem is to answer a question about the value of knowledge, namely: In what way (or how, or why, or in what sense) is knowledge more valuable than (mere) true belief?<sup>1</sup> This is the question that contemporary epistemologists discuss, when they discuss the value problem. Moreover, they make two assumptions about this question. First, they understand the question as presupposing that knowledge *is* more valuable than true belief. Second, they typically assume that true belief is valuable. This is perhaps because we often motivate the problem by appeal to a point from Plato’s *Meno* (97b): that from a practical point of view knowledge seems no more valuable than true belief. The person who knows the way to Larissa seems in no better a position, as far as her project of traveling to Larissa (which is ‘practical’, just in the sense that it involves her performance of certain actions, namely, those that comprise traveling to Larissa), than someone who has a true belief, not amounting to knowledge, about the way to Larissa. What explains, then, the fact that knowledge is more valuable than true belief?

True belief, not amounting to knowledge, clearly has some practical value (indeed, the problem is motivated by the fact that it seems to have all the practical value one could want in a belief). So when we motivate the *Meno* problem in this way, we introduce the assumption: that true belief is valuable. However, one could respond to the value problem, as posed in the *Meno*, by first conceding that knowledge is no more *practically* valuable than true belief, and then going on to argue that the value of knowledge is of some *non-practical* variety. Socrates, for his part, seems not to opt for this route: he argues that knowledge has a kind of stability and security, compared to true belief, and that it is therefore of greater practical value. Consider Duncan Pritchard’s example: “if one knows the way to Larissa ... then one is less likely to be perturbed by the fact that the road, initially at least, seems to be going in the wrong direction.” (2007)

---

[Standard “Greatness of Mind” paper acknowledgements.]

<sup>1</sup> In what follows I’ll drop the ‘mere’. Also, here I treat ‘having value’ and ‘being valuable’ as synonymous.

However, many of those who chose to answer the value problem by appeal to some non-practical species of value typically make the *Meno*-esque assumption that true belief has *some* of that species of value. So among those who take knowledge to have distinctive *epistemic* value, many state the value problem as a question about what additional epistemic value knowledge has, in addition to the epistemic value of true belief. It is this way of conceiving of the question that gives rise to the “swamping problem” for externalist theories of justification (Kvanvig 2003, pp. 44-51); true belief must be assumed to have some epistemic value, if said value is going to “swamp” the value of knowledge.

It is possible to solve the value problem straightaway by denying the epistemic value of true belief, so long as a case can be made that knowledge enjoys a modicum of epistemic value. But many would consider this a drastic move. Most epistemologists want to solve the value problem without denying that true belief has some epistemic value. For them, in any event, the value problem therefore involves answering a question about what additional epistemic value knowledge has, in addition to the epistemic value of true belief.

### *1.2 The sources problem*

What does it mean to say that true belief has epistemic value? There are two related questions here. The first concerns what it means to say of something valuable that its value is ‘epistemic’. This question has received remarkably little attention, given the ubiquity of the term ‘epistemic’ in contemporary philosophy. It is sometimes used as a synonym for ‘doxastic’, which means something like ‘having to do with belief’. So, for example, some speak of ‘epistemic agents’, meaning people, in their capacity as believers. (There is an instructive analogy here with certain uses of the word ‘moral’: consider ‘moral psychology’, which concerns a wide range of topics pertaining to thought, feeling, and action, but about which there is nothing especially moral, in the contemporary sense of that word.) The ‘epistemic’ is often contrasted with the practical, suggesting synonymy with ‘theoretical’. I’ll return to the issue of defining ‘epistemic’ in a moment.

The second question involved, in asking what it means to say that true belief – or anything else – has epistemic value, is a question about the source or sources of normativity, and in particular about the source or sources of normativity involved in the particular kind of evaluation in question, in this case epistemic evaluation.<sup>2</sup> What do I mean by this?

There are a plurality of practices that could each plausibly be called a type of epistemic evaluation. Any such practice will consist of two things: (i) a division or classification of things, and (ii) a hierarchy (i.e. positive and negative valences) applied to the elements of the division. Thus we can evaluate beliefs as rational (positive) or irrational (negative), as amounting to knowledge (positive) or not (negative), as justified (positive) or unjustified (negative), and finally as true (positive) or false (negative). This is not an exhaustive catalog of varieties of epistemic evaluation, but these are paradigmatic types of epistemic evaluation.

---

<sup>2</sup> Cf., obviously, Korsgaard 1996. The sources problem that I identify here is, on its face, quite different from the “normative question” posed by Korsgaard. I think that, at heart, these treat of the same issue, or at least very closely related issues, but I don’t assume that here.

Furthermore, these evaluations are *normative*. This is what is meant by saying that epistemic evaluation partly consists of a hierarchy applied to the classified elements. Evaluation is not mere classification, but classification colored by normativity. But what does this amount to?

One way of seeing the normativity involved here is to notice that epistemic evaluation, at least in some cases, involves what Strawson called “reactive attitudes.” There is, at least sometimes, a kind of blame involved in evaluating someone’s belief that  $p$  as irrational; there is, at least sometimes, a kind of praise involved in saying that someone has knowledge of a certain proposition. Irrational believing can inspire our resentment and disapprobation; epistemically competent cognition can inspire our admiration and respect. This reveals, further, that epistemic evaluation, at least sometimes, is the evaluation of persons – to evaluate a belief as irrational is, at least sometimes, to evaluate a certain believer as having fallen short of some standard; to evaluate a belief as an instance of knowledge is, at least sometimes, to evaluate a certain knower as having done well, epistemically. (Cf. §1.4 and §2.3)

Given that epistemic evaluation is normative, we can ask a question about the ‘source’ of this normativity. What justifies, or grounds, or warrants, or legitimates (i.e. explains the legitimacy of), this form of evaluation?

What does this question mean? To see what it means, consider a form of intuitively illegitimate evaluation of beliefs (i.e. doxastic evaluation), which evaluates positively those formed on a Tuesday, and which evaluates negatively those formed on any other day of the week. On this scheme of evaluation, people who form beliefs on Tuesdays are, at least sometimes, praised in virtue of that fact, and people are thought of disapprovingly, at least sometimes, for having beliefs not formed on a Tuesday. The fundamental norm, we might say, for this species of doxastic evaluation, is:

TUESDAY NORM: S’s belief that  $p$  is correct iff S formed her belief that  $p$  on a Tuesday.

‘Correct’ here should be taken as a generic term of approbation; all that is meant is that beliefs formed on Tuesdays are taken to be better, on this scheme of evaluation, than beliefs formed on any other day. Tuesday beliefs have the positive valence in this scheme of evaluation; other beliefs have the negative valence.

We want to say, I think, that evaluation vis-à-vis the Tuesday norm is somehow illegitimate. Or, at least, that epistemic evaluation of beliefs is legitimate in a way that evaluation vis-à-vis the Tuesday norm isn’t. Consider simply the evaluation of beliefs in terms of their truth or falsity. Epistemic evaluation involves, at least, the evaluation of beliefs vis-à-vis the following:

TRUTH NORM: S’s belief that  $p$  is correct iff  $p$  is true.<sup>3,4</sup>

---

<sup>3</sup> Note that this rule doesn’t just derive from the intrinsic value of truth (if any such there be, and whatever exactly such value would amount to): if I imagine the truth, my imagining

What makes it the case (and in what sense is it the case) that epistemic evaluation of beliefs (e.g. doxastic evaluation vis-à-vis the truth norm) is legitimate, while evaluation of beliefs vis-à-vis the Tuesday norm is not? Answering that question is the *problem of the sources of epistemic normativity*, or the *sources problem*, for short. A solution to the sources problem will answer both of our questions, above. It will define (or characterize or articulate) the notion of the ‘epistemic’, and it will explain the legitimacy of epistemic evaluation.

### 1.3 Three responses to the sources problem

Here I’ll describe three ways of responding to the sources problem: epistemic essentialism, epistemic conventionalism, and pragmatism about epistemic evaluation. Then in the second part of the paper, I’ll criticize one species of epistemic essentialism.

The word ‘epistemic’ is sometimes taken to be synonymous with the phrase ‘purely intellectual’, so that ‘epistemic evaluation’ is taken to be synonymous with ‘purely intellectual evaluation’. Such evaluation is to be contrasted with various forms of non-epistemic evaluation of beliefs: my believing that  $p$  might be conducive to my wellbeing, and thereby do well vis-à-vis prudential evaluation, or my belief’s content might be particularly beautiful (an elegant mathematical theorem, for example), and thereby do well vis-à-vis aesthetic evaluation. But neither of these evaluations will be ‘purely intellectual’. Evaluation vis-à-vis the Tuesday norm, by the same token, is not ‘purely intellectual’ evaluation. Evaluation vis-à-vis the truth norm, by contrast, is ‘purely intellectual’ evaluation. This idea goes back to Aristotle, who wrote that:

[O]f the intellect which is contemplative, not practical nor productive, the good and the bad state are truth and falsity respectively (for this is the work of everything intellectual). (*Nicomachean Ethics*, 1139a27-9)

But this just pushes the question back a step. For now we must ask after the legitimacy of ‘purely intellectual’ evaluation. But we now have a clue as to how to proceed. Perhaps what it means to say that epistemic evaluation is ‘purely intellectual’ is that epistemic evaluation bears some special relation to the intellect, in other words, to the mind, to cognition, or to believing, as such. Epistemic evaluation of beliefs, so the thought goes, is the evaluation of beliefs, as such, in other words, the evaluation of beliefs, qua beliefs. More generally, we could say that the epistemic evaluation of any act of intellection, as it were, is the evaluation of that act of intellection, qua act of intellection. (Believing, withholding, inquiring, and so on, are ‘acts of intellection’, in the intended sense, and by calling them ‘acts’ we do not yet mean to imply anything about whether they are voluntary, etc.)

---

doesn’t merit positive epistemic evaluation, at least not in the way that a true belief merits positive epistemic evaluation.

<sup>4</sup>To say that epistemic evaluation involves evaluation vis-à-vis the truth norm leaves open the possibility that there are other fundamental epistemic norms. So what I have said so far is neutral on the question of whether truth is the only fundamental epistemic value or good.

The idea here is that, for example, to evaluate someone's beliefs vis-à-vis the truth norm is to evaluate them as the kinds of things that they are, whereas to evaluate someone's beliefs vis-à-vis the Tuesday norm is to apply some alien standard to said beliefs. In other words:

EPISTEMIC ESSENTIALISM (ABOUT DOXASTIC EVALUATION): Epistemic evaluation of beliefs is the evaluation of beliefs, qua beliefs, and this explains why the epistemic evaluation of beliefs is legitimate.

What it means to say, of a species of evaluation, that it is the evaluation of "beliefs, qua beliefs" remains to be explained (§1.4). Setting that aside, why would the fact that epistemic evaluation is the evaluation of beliefs, qua beliefs, explain why epistemic evaluation of beliefs is legitimate? The idea is that the essence or nature of belief, as such, exposes it to epistemic evaluation, as opposed to other forms of doxastic evaluation (such as evaluation vis-à-vis the Tuesday norm). A knife, for example, in virtue of being the kind of thing that it is, in virtue of its nature or essence, can appropriately be evaluated for sharpness, strength, weight, and other such features that intuitively go into the evaluation of a knife, qua knife. To complain of a knife that it does not produce beautiful music, for example, is to apply an alien standard to it, as opposed to a standard that flows from the nature of essence of knives, as such. (More on the evaluation of artifacts in §2.3.)

In this paper I am going to focus on species of epistemic essentialism that are based on the following idea:

TRUTH NORM ESSENTIALISM (ABOUT DOXASTIC EVALUATION): The evaluation of beliefs vis-à-vis the truth norm is (at least one form of) the evaluation of beliefs, qua beliefs.

This is sometimes put by saying that truth is a "constitutive standard of correctness" for belief, or that truth is the "aim of belief."<sup>5</sup> By contrast, being formed on Tuesday is not a constitutive standard of correctness for belief. The truth norm flows from the nature or essence of belief, whereas the Tuesday norm does not. In what follows I'll call the conjunction of epistemic essentialism and truth norm essentialism simply 'essentialism'.

The essentialist needs to explain how the essential applicability of the truth norm to belief explains the legitimacy of epistemic evaluation of beliefs, in general. She needs to provide an account that takes us from the idea that the evaluation of beliefs vis-à-vis the truth norm is the evaluation of beliefs, qua beliefs, to the idea that epistemic evaluation of beliefs, in general, is the evaluation of beliefs, qua beliefs. Let's assume, for the sake of argument, that this account can be given.

In §2 I criticize three versions of one species of essentialism. What are the alternatives to essentialism?

---

<sup>5</sup> See Williams 1976, Humberstone 1992, Railton 1994, Velleman 2000, Wedgwood 2002, Shah 2003, Boghossian 2003 and 2005, Shah and Velleman 2005, Steglich-Petersen 2006, and Sosa 2009.

One response to the sources problem is to reject the presupposition that epistemic evaluation of beliefs is justified or legitimate in some way that prudential evaluation, or aesthetic evaluation, or evaluation vis-à-vis the Tuesday norm, is not. Evaluation vis-à-vis the Tuesday norm is, of course, non-epistemic evaluation, as are prudential and aesthetic evaluation of beliefs, but this means nothing more than that these are different species of doxastic evaluation. None has primacy or precedence against any of the others. Species of evaluation can be individuated by the norms or standards or values, vis-à-vis which things are evaluated, and epistemic evaluation of beliefs is nothing more than the evaluation of beliefs relative to a certain set of norms or standards or values.

EPISTEMIC CONVENTIONALISM: Epistemic evaluation is evaluation vis-à-vis certain norms or standards or values, and has no particular or special legitimacy, when it comes to the evaluation of beliefs.

An obvious candidate for a fundamental epistemic value is, of course, truth. (There may be other fundamental epistemic values, and at least some derivative epistemic values could be accounted for in terms of truth.)<sup>6</sup> We can characterize epistemic evaluation as evaluation relative to the “goal” of attaining the truth and avoiding error.<sup>7</sup> But there is nothing further to be said about why such evaluation is legitimate, when it comes to beliefs. Epistemic evaluation of beliefs is not claimed to be any more or any less legitimate than prudential or aesthetic evaluation of beliefs, it is simply one among a plurality of species of doxastic evaluation.

Both essentialism and conventionalism (in this ‘truth goal’ version) posit truth as a fundamental epistemic value. The difference comes from how the theories explain this fact. The essentialist explains the epistemic value of truth by appeal to the nature of belief. The conventionalist explains the epistemic value of truth by appeal to the nature of the epistemic, where this is simply a conventionally demarcated domain of evaluation.

Essentialism explains the legitimacy of epistemic evaluation by appeal to the normativity of belief. Conventionalism says that the legitimacy of epistemic evaluation can’t be explained, and that it is therefore no more legitimate than any other species of evaluation. *Pragmatism about epistemic evaluation*, by contrast, explains the legitimacy of epistemic evaluation in terms of our non-epistemic interests and concerns. Our concern for survival, for example, might justify a quite general evaluative practice, which takes true and reliably formed beliefs to be better than false or unreliably formed ones – i.e. it might justify the practice of epistemic evaluation. In answer to the question of why epistemic evaluation is legitimate, the pragmatist answers that beliefs that do well epistemically are also, typically, in general, or for the most part, beliefs conducive to the attainment of our various non-epistemic interests

---

<sup>6</sup> I favor a conception of the conventional meaning of ‘epistemic’ that emphasizes the idea that epistemic evaluation prizes impartial and unbiased pursuit of the truth. I therefore concur with Richard Feldman’s characterization: “Epistemic obligation concerns obligations to believe to which the practical benefits of beliefs are not relevant. They are obligations that arise from a purely impartial and disinterested perspective.” (1988, p. 236)

<sup>7</sup> On the truth goal, see Alston 1985, Bonjour 1985, pp. 5-8, Goldman 1987, p. 98, Foley 1987, p. 125, David 2001, Sosa 2001, and Zagzebski 2003. But none of these theorists is obviously a conventionalist, in my sense.

and concerns. Epistemic evaluation is there a species of evaluation vis-à-vis those various non-epistemic interests and concerns.

#### 1.4 Three species of essentialism

There are three species of essentialism: biological essentialism, normative essentialism, and performance essentialism. Let us consider each in terms of how it explains the legitimacy of applying the truth norm to beliefs.

*Biological essentialism* maintains, roughly, that being true is the biological, which is to say evolutionary, function of belief. The analogy here is to bodily organs, for example, the heart, whose biological function is to pump blood. A heart that pumps blood is doing well, qua heart, because such a heart is functioning properly, qua heart. By the same token, a belief that is true is doing well, qua belief, because such a belief is functioning properly, qua belief. The truth norm could be said to codify this fact.<sup>8</sup> I won't discuss biological essentialism in this paper. I will mention only that biological essentialism makes the legitimacy of epistemic evaluation an a posteriori matter. Just as the question of the heart's biological function is an empirical matter, so is the question of belief's biological function. This species of essentialism cannot be established a priori.

Biological essentialism seeks to reduce epistemic normativity to the normativity of biological functions. By contrast, *normative essentialism* maintains that epistemic normativity is irreducible. On this view, the truth norm is a necessary conceptual truth about belief. Truth is the "constitutive standard of correctness" for beliefs; to employ the concept of *belief* is to employ a normative concept, which includes this constitutive standard.<sup>9</sup> But this normativity is not to be reduced or explained away (as on biological essentialism, and as on performance essentialism, to be articulated in a moment). On this view, epistemic normativity is the normativity of belief, as such; in that sense it is *sui generis* (by contrast, again, with biological and performance essentialism). I won't discuss normative essentialism here, but note that it seems incompatible with a certain kind of naturalism.

By contrast, naturalists should be happy with accounts of normativity that explain it away in terms of the desires, intentions, or goals of individuals, i.e. in terms of the literal aims of individuals. It is one thing to say that belief aims at truth, and that this is an irreducible normative fact built into the concept of belief. It is another to say that *believers* literally aim at truth: that they desire it, intend to get it, have it as their goal, or whatever. On the assumption – which is a substantial, but not unreasonable assumption – that a naturalistic account of literal aiming (desiring, intending, having a goal) can be given, the theory that explains epistemic normativity in terms of literal aiming will thereby achieve a completely satisfying naturalistic reduction of epistemic normativity. Let me explain with some specifics, which would be different on different versions of the species of essentialism I'm describing.

---

<sup>8</sup> Cf. Millikan 1987, pp. 93-4. See also Dennett 1971, pp. 101-3.

<sup>9</sup> See Dennett 1971, Davidson 1983, Wedgewood 2002 and 2007, and Boghossian 2003 and 2005. Shah 2003. Shah and Velleman 2005.

Our question concerned the legitimacy of evaluating beliefs vis-à-vis the truth norm, i.e. the legitimacy of evaluating them as correct when they are true, and incorrect when they are false. Suppose that, necessarily, anyone who believes  $p$  also desires to abide by the truth norm – that, at a minimum, she desires to believe  $p$  only if  $p$  is true.<sup>10</sup> Someone who believes  $p$ , where  $p$  is false, therefore, will have failed to attain something that she desired. Her belief will be a failure, by her own lights, in terms of what she wanted. This, so the account goes, is the sense in which beliefs are correct if and only if true: believers fail to get what they want, when they believe falsely, and succeed in getting what they want, when they believe truly.

Any view that, in similar fashion, explains the legitimacy of epistemic evaluation of beliefs in terms of the desires, intentions, or goals of believers is a version of *performance essentialism*.

Performance essentialism treats epistemic normativity as a special case of *performance normativity*. Essential to such normativity is the idea that a performance fails if and only if the performer (i.e. the person that performs) fails. Indeed, something stronger: the reason that a particular performance amounts to a failure is that the person who performs fails to do something that she desires to do, intends to do, or that it is her goal to do. My performance is a failure, in such cases, because *I* am a failure. For performance normativity exists *only when* I try to do something, and fail. The ‘source’ of normativity, in such cases, is my own intending or desiring or goal-setting. For this reason, the failure in cases of performance normativity is always failure *by my own lights*. When I try to  $\Phi$ , being an instance of  $\Phi$ ing becomes a standard of correctness for my performance, but *only* because I *set that standard for myself*.<sup>11</sup>

## 2. Performance essentialism: three varieties

I am going to criticize performance essentialism, by criticizing the three most plausible ways of developing it.

### 2.1 Believers vs. archers

---

<sup>10</sup> Would she also have to desire to believe  $p$  if  $p$  is true? I think not. The truth norm codifies a standard of correctness for believing: it tells us what a belief would have to be like, if it’s going to be a correct belief. The norm is therefore conditional. Compare: Baron Points are the only oysters worth eating, which would normally be taken to mean that an oyster is worth eating iff it’s a Baron Point. You can satisfy this norm without eating every Baron Point, indeed you can satisfy it without eating Baron Points at all. You just can’t eat oysters that aren’t Baron Points. (That an oyster is worth eating doesn’t mean you fail, by not eating it – for you may have no interest in eating, i.e. you may not be hungry.) For the same point, see Sosa 2001.

<sup>11</sup> We can individuate performances in terms of their literal aims. When I try to make a cornbread, part of what it is to be doing what I am doing is to be trying to make a cornbread. I would not be doing what I am doing were I not trying to make a cornbread. If this is right, then the standards of correctness involved in performance normativity will be constitutive standards of correctness.

The performance essentialist might try to explain the legitimacy of the epistemic evaluation of beliefs by appeal to the believer's curiosity. The curious believer wants to know the truth; if she winds up with a false belief, then she has failed by her own lights. But, as Stephen Grimm (2008) has shown, curiosity does not provide an adequate explanation of the legitimacy of epistemic evaluation. For we are selectively curious, and there are many things about which we haven't the slightest curiosity:

Suppose ... that the endless wait to see the doctor has driven you to count the number of motes of dust on the doctor's waiting-room coffee table. After a few minutes, you conclude that there are 53 motes. (2008, p. 742)

Your motive for counting the motes is not curiosity about their number, but sheer boredom; you count them not because you are curious about their number, but because you are bored.

For all that, however, we evidently *can* evaluate – and in certain cases *would* evaluate – the epistemic status of your belief. (Ibid.)

The legitimacy of epistemic evaluation is therefore not explained by appeal to the curiosity of believers, and the reason for this is that beliefs, which are intuitively subject to epistemic evaluation, are often formed absent curiosity. Indeed, the belief-forming practices of the incurious are often precisely those which we want to evaluate *negatively* from the epistemic point of view.

However, isn't there some sense in which you *are* curious about the number of motes, at least once you begin counting them, or at least once you get to the point where you are forming a belief about their number? Well, again, not in the ordinary sense: if asked whether you had ever been curious about the number of motes, you'd reply that you were never curious, but merely bored to death. You could, of course, be in bad faith about your curiosity – embarrassed at being curious about something so trivial, at least by other people's lights. But the point here is that you easily could be absolutely authentic in your insistence that curiosity about the motes was never present. Which just means that, if there is a sense in which you were curious, it is not the ordinary sense.

What we seek is a sense in which *every* belief involves curiosity, for epistemic evaluation intuitively applies to beliefs, in general, not just those whose formation was partially the result of curiosity, or any other such motivation involved in some cases of believing, but not in others. As Richard Feldman points out, epistemic requirements apply to everyone, regardless of idiosyncratic motivations or desires. Evidentialism, for example, "says that all people epistemically ought to follow their evidence, not just those who have adopted some specifically epistemic goals." (2000, p. 682) What Grimm and Feldman are on to is:

APPLICATION REQUIREMENT (ON SOLUTIONS TO THE SOURCES PROBLEM):  
Solutions to the sources problem must explain the legitimacy of epistemic evaluation in *all* cases of believing.

With this in mind, consider the following claim from Lloyd Humberstone:

[U]nless one takes there to be a criterion of success in the case of an attitude towards the proposition that  $p$ , and, further, unless that criterion is truth, then whatever else it may be, the attitude in question is not that of belief. So unless the attitude-holder has what we might call a controlling background intention that his or her attitudinizing is successful only if its propositional content is true, then the attitude taken is not that of belief. (1992, p. 73)

The idea must be that, necessarily, someone believes  $p$  only if she intends that she believe  $p$  only if  $p$ . This satisfies the Application Requirement.<sup>12</sup> But it seems manifestly implausible. Consider Obama, who believes that he is in Washington. (It's awkward to *say* that Obama believes that he is in Washington, because this implies that he isn't; nevertheless, he does believe that he's in Washington.) But Obama does *not* intend that he believe he's in Washington only if he is in Washington (or anything like that), because he does not intend anything with respect of this particular belief. Obama's believing that he is in Washington is probably something that he has no thoughts about whatsoever, and therefore with respect of which he never forms any intentions.<sup>13</sup>

The problem is not, I think, in the specification of the required intentional content. The problem is not going to be solved by, for example, requiring a more general intention to, for example, believe  $q$ , for all  $q$ , only if  $q$ .<sup>14</sup> Even if reflective believers like Obama have such intentions, many believers do not: consider human children and non-human animals. If children and animals do not have such intentions – and I think it would be highly implausible to say that they do – then Humberstone's claim entails that they do not believe anything. But this, too, is highly implausible.<sup>15</sup> The point generalizes: people who believe

---

<sup>12</sup> The formulation is virtuous for two other reasons. First, it's better than saying that believers must desire that their beliefs be true, which ambiguously suggests that no one could believe  $p$  while wishing  $p$  were false. Second, the requirement doesn't run afoul of the considerations discussed in footnote 10. For an alternative account that appeals to the desire for truth, see Steglich-Petersen 2009.

<sup>13</sup> See Owens 2003 for additional problems for performance essentialisms that conceive of the aim of belief in intentional terms.

<sup>14</sup> Thanks to Crispin Wright for the suggestion.

<sup>15</sup> Two points. First, my philosophical interest is in the concept of belief on which knowledge is a species of belief (or, if you like, on which knowing entails believing). It is, I think, even more clear that children and animals *know* various things. Their behavior would be inexplicable otherwise. Bernard Williams (1973) suggests that knowledge is not a species of belief, on the way to arguing that belief aims at truth. Clearly, if the thesis that belief aims at truth is to provide any help for the sources problem, it will have to be understood as a thesis about the concept of belief on which knowledge is a species of belief. (It's beyond the scope of this paper to consider the view, defended by Davidson (1982), that belief requires second-order thought, and therefore that animals do not believe anything.) Second, the question of animal intelligence has always been at the heart of the question of normativity. Kant and neo-Kantians posit (irreducibly normative) categorical imperatives that govern belief and action, as such. Hume and neo-Humeans reject categorical imperatives, countenancing only the hypothetical variety. It is not a coincidence that Humeans extend thought and action to non-human animals, taking these to be common natural phenomena, while Kantians take belief and action to be rare phenomena, and rather hard to come by.

that  $p$  do not normally intend or desire or have as their goal anything about their beliefs that  $p$ ; they may have no second-order thoughts whatsoever about their beliefs. Thus we should reject Humberstone's claim, and any variation on it that requires truth-directed intentions (e.g. to believe  $p$  only if  $p$ ), or desires or goals, to be had by all believers.

If this is right, performance essentialism, which seeks to explain the legitimacy of the epistemic evaluation of beliefs by appeal to the desires, intentions, or goals of believers, faces a dilemma:

DILEMMA: Either (I) the relevant desires, intentions, or goals are idiosyncratic, i.e. they are desires, intentions, or goals that some believers have, and some believers don't have, which runs afoul of the Application Requirement, or else (II) the relevant desires, intentions, or goals are essential to believing, which is implausible: no desires, intentions, or goals seem necessarily required for believing.<sup>16</sup>

Is there a way to save performance essentialism? Ernest Sosa has recently proposed a picture of belief that seems designed to avoid our dilemma.<sup>17</sup> Sosa compares beliefs with the shots attempted by an archer, which shots are evaluable in terms of whether they succeed in their aim, namely, hitting the target. (2007, p. 22) Such shots are performances, and in general:

A performance is better than otherwise for not having *failed*, i.e., for not having fallen short of its objective. In line with that, it is *good* if it succeeds, if it reaches its objective. A performance is at least good *as such* for succeeding. (2009, p. 9)

Believing is a performance whose objective is truth; thus a belief is good, as such, if true. Elsewhere, Sosa sums up this account of epistemic normativity:

The good shot is the central value that organizes the sport of archery and the criticism proper to it. [...] All such evaluation is dependent on the value constituted by the good shot. [...] Truth is similarly the fundamental value of epistemology. (2007, p. 77-8)

But recall the application requirement. Unless the archer intends to hit the target, her shot (or perhaps her "shot") cannot be normatively evaluated as successful if and only if it hits

---

<sup>16</sup> Cf. Grimm 2009, pp. 249-51.

<sup>17</sup> See Sosa 2007, pp. 22-3 and pp. 77-91, and 2009. It should be noted that Sosa does not propose his view as a solution to the sources problem, but rather as a solution to the value problem. The version of performance essentialism I attribute to him seems implicit in some recent work on the value problem; I have in mind those theorists whose solution to the value problems involves the idea that knowledge is a credit-worthy achievement (Zagzebski 1996, Riggs 2002, Greco 2003). These views solve the value problem by appeal to the idea that people who know are responsible for believing the truth, by contrast with those who stumble upon it. In as much as this picture suggests that believing is a goal-directed activity, it suggests a version of performance essentialism.

the target. If she releases an arrow by mistake, for example, not intending to shoot, then her shot cannot be evaluated as successful or unsuccessful. Alternatively, if she shoots at something other than the target – if she aims to shoot a rival, who foolishly chose to stand close to the archery range – then her shot is successful if and only if it strikes the rival, and *not* if and only if it strikes the target.<sup>18</sup> We can see now that (i) the normativity of archery depends, essentially, on the intentions (or desires, or goals) of the archer, and (ii) there is no aim that is constitutive of archer's shots – archers can aim at various things (targets, rivals, etc.), and the evaluation of a particular shot will be relative to the aim, in taking that shot, of the particular archer who took it. However, Sosa seems aware of this difficulty, and this is where it seems his view might avoid our dilemma. Here is what he says:

Some acts are performances, of course, but so are some sustained states. Think of those live motionless statues that one sees at tourist sites. Such performances can linger, and need not be constantly sustained through renewed conscious intentions. The performer's mind could wander, with little effect on the continuation or quality of the performance. Beliefs too might thus count as performances, long-sustained ones, with no more conscious or intentional an aim than that of a heartbeat. (2007, p. 23)

Is belief such a performance? Sosa offers two analogies for belief in this passage. First, he compares believing to a mime's performance, in which the mime stands motionless over some interval of time, during which she does not consciously intend to perform. Second, he compares believing to the beating of a person's heart, which seems a kind of performance, although it involves no intentions (whether conscious or not) on the part of the person in question. Is believing analogous to either of these kinds of performance?

The answer is No. The mime's performance, although it remains a performance through the interval (during which the mime no longer consciously intends to perform), is unlike believing in two relevant respects. First, Sosa clearly imagines that the mime initiates her performance with an intention to perform; thus the idea that she need not "renew" her intention to perform. But an episode of believing seems to require neither an intention (e.g. to believe  $p$  only if  $p$ ) that is present throughout the episode, *nor* an initiating intention, present perhaps only at the beginning of the episode. Second, while it sounds fine to say that a mime may perform, at  $t$ , while not, at  $t$ , consciously intending to perform, it sounds much less plausible to say that a mime may perform, at  $t$ , while not, at  $t$ , intending to perform. When the mime's mind wanders, during such a performance, she no longer consciously intends to perform. But does she no longer intend to perform? Intentions, like beliefs and desires, need not be conscious.

Recall the point above, about Obama. The point there was not merely that attributing to Obama a *conscious* intention (to believe  $p$  only if  $p$ ) would be implausible, although that seems particularly implausible. The point was that second-order thoughts, about one's own beliefs, are rare. They are certainly not required to have beliefs in the first place.

---

<sup>18</sup> It could be objected that this is not a genuine shot, in the sense that it is not a shot in an archery competition. I take up the analogy between believing and game-playing in §2.2.

What about Sosa's example of the heartbeat? Here I think we have stretched the notion of a 'performance' too far. Consider the point that Sosa is trying to make, that the heart beats, regardless of the intentions of the person within whom it beats, and consider the language needed to make it. It is *the heart*, and not the person, that beats; that's why we say that it beats within her. She does not beat her heart, as it were, for the beating of one's heart is not something that one *does*. This just shows that it is not a performance, in the relevant sense. But perhaps the heartbeat was simply a poor choice of example. We would say that I digest my food, and this is also something that happens without my intending it to happen. The real issue here is that, once we analogize belief to entirely unintentional activities, we have abandoned performance essentialism. The explanation of epistemic evaluation, by comparing belief to biological functions (digestion) or organs (hearts) is (what I called above) biological essentialism. Sosa's picture, with its leading analogy of the archer, does not seem to be such a view.

One final objection. Might we say that believers always have a truth-directed intention, but that this intention is tacit, implicit, or dispositional? This depends on what it means to say that some mental state is 'tacit', 'implicit', or 'dispositional'. One way of understanding these notions is as follows:

S is *tacitly*<sub>1</sub> in mental state M iff (i) S is in M, and (ii) S has never consciously considered the fact that she is in M. Such a state can be said to be *tacit*<sub>1</sub>.<sup>19</sup>

I think that most of our mental states are *tacit*<sub>1</sub>. And it is not plausible that Obama *tacitly*<sub>1</sub> intends to believe that he is in Washington, only if he is in Washington. Again, the objection is not that Obama's truth-directed intention isn't conscious; it's that it's implausible to say that he has such an intention at all.

Alternatively, consider:

S is *tacitly*<sub>2</sub> in mental state M iff (i) S is not in M, but (ii) S would be in M, were she in appropriate circumstances. Such a state can be said to be *tacit*<sub>2</sub>. For the mental state of belief, for example, S *tacitly*<sub>2</sub> believes *p* iff (i) S doesn't believe *p*, but (ii) S would believe *p* were she to consider the question of whether *p* is true.

Unlike a *tacit*<sub>1</sub> belief, which is a kind of belief, a *tacit*<sub>2</sub> belief is no more a belief than an egg is a chicken. It is not obvious what the "appropriate circumstances" would be to test for someone's *tacit*<sub>2</sub> intentions; perhaps *tacitly*<sub>2</sub> intending to  $\Phi$  requires that one intend to  $\Phi$ , upon considering the question of whether to  $\Phi$ . But setting that aside, if one were to claim that Obama *tacitly*<sub>2</sub> intends to believe he's in Washington only if he is in Washington, it would be no objection to this claim to say that Obama intends no such thing. So the issue before us is whether believing requires a truth-directed intention, which may be *tacit*<sub>2</sub>.

It seems possible that it does. Adopt the assumption just mentioned, that S *tacitly*<sub>2</sub> intends to  $\Phi$  iff (i) S doesn't intend to  $\Phi$ , but (ii) S would intend to  $\Phi$  were she to consider the

---

<sup>19</sup> Cf. Churchland 1989, p. 390.

question of whether to  $\Phi$ . Could there be someone who believes that  $p$ , and yet upon considering the question of whether to believe  $p$  only if  $p$ , would not intend to believe  $p$  only if  $p$ ? If we had taken someone to believe something, we might reconsider the attribution of belief to her, upon discovering that she cares not a whit for the truth of her supposed “belief.” So perhaps believing does require a “controlling background intention,” but one that may be merely tacit<sub>2</sub>. We might have to come up with some fancy “appropriate conditions” to coax the right intentions out of children and animals, but perhaps this could be done.

The problem here is trying to account for the legitimacy of epistemic evaluation. Recall Sosa’s idea: a performance that doesn’t reach its aim is a failure, thus a successful performance is good, as such. True belief is therefore good, as such, i.e. as an instance of a performance that achieves its aim. Let’s translate this into the language of intentions: when someone intends to  $\Phi$ , but doesn’t  $\Phi$ , it’s a failure, thus when someone intends to  $\Phi$ , and does  $\Phi$ , she has done something good, qua intentional performance. But this is not true of tacit<sub>2</sub> intentions. It’s no failure if I don’t achieve what I tacitly<sub>2</sub> intend to achieve. There are plenty of things that I *would* intend to do, *were* I to think about whether to do them. My not doing them does not amount to a *failure*, of the sort that would warrant negative evaluation.

I stroll down Fifth Avenue, without a care in the world. The following might be true of me: were I to consider whether to check out the cufflinks at Zegna, I would form the intention to check them out. This isn’t because I’ve been meaning to check them out – i.e. because I had already formed the intention to check them out, at some earlier time – it’s simply that the thought has never crossed my mind. And as it happens, the thought still doesn’t cross my mind, even as I walk by Zegna without going in.

Was this a failure? Did something bad happen? Among the thousand natural shocks that flesh is heir to, we will not find the tragedy of our not getting what we merely would have wanted.

Thus, it is implausible that believing requires a tacit<sub>1</sub> truth-directed intention, and not normatively significant if and when believing involves a tacit<sub>2</sub> truth-directed intention.

## 2.2 *Believers vs. game players*

On Sosa’s version of performance essentialism, a belief is analagous to a shot. As the shot aims at its target, so belief aims at truth. My objection to this can be distilled into the following observation: an archer’s shot aims at a target (i.e. is evaluable as successful vis-à-vis the standard of hitting that target) *only* in as much as the *archer* literally aims at that target. Once that is recognized, it is not plausible to analogize belief to the archer’s shot, for believers need not literally aim at truth.

However, it could be objected that I have misunderstood the analogy. I considered archery as a special case of intentional activity, in which an individual aims for some goal (the archer intends to his some target), and either succeeds or does well (e.g. by hitting the target) or fails or does poorly (e.g. by not hitting it). Imagine, however, an archery *competition*, and imagine (as above) the archer who releases an arrow by mistake. Above I said that such a

“shot” could not be evaluated vis-à-vis the standard of hitting the target. But this seems wrong in the context of an archery competition: if the archer releases an arrow by mistake, but when it happens to be her turn to shoot in the competition, her shot might still count as a successful shot, i.e. one that counts towards her total score in the competition, or one that wins her the praise of the cognoscenti, or whatever.

This suggests an analogy between belief and archery, where archery is understood not as the intentional action of an individual trying to hit a target, but where archery is understood as a *game*, in which hitting the target is the constitutive standard of correctness or success. An archer who doesn’t intend to hit the target, so the argument would go, is no less evaluable for her lack of intention; her shot is a failure, *qua* shot. Analogously, a believer whose belief is false is no less evaluable for her lack of a truth-directed intention; her belief is a failure, *qua* belief.

A performance essentialism of this kind is suggested by Ralph Wedgwood’s comparison of the truth norm to the rules of chess.<sup>20</sup> Just as it violates the rules of chess to make an illegal move, it violates the rules of “reasoning” (which is the practice that includes “forming and abandoning beliefs”) to form a false belief (and just as it would be irrational to make a move one knows to be illegal, it would be irrational to form a belief one knows to be false). (2002, pp. 268-9) The rules of chess determine the constitutive standard of correctness for the ordinary practice of chess playing; the rules of reasoning determine the constitutive standard of correctness for the ordinary practice of reasoning. As above, what makes a false belief an incorrect “move” in the game of reasoning doesn’t depend on the believer’s having intended to believe truly, just as the illegality of a move in chess doesn’t depend on the believer’s having intended to make a legal move: someone might in no uncertain terms intend to make an *illegal* move, and the move would still be *wrong*, given the rules of the game. If the analogy with believing sticks, the same could be true of the wrongness of false beliefs.

But the analogy does not stick. The important reason for this is that playing a game – being part of an archery competition, playing chess – necessarily requires certain intentions, analogues of which are not necessarily present in believing. We can evaluate the archer’s unintentional shot *only* because the archer is entered in the competition, and she can only be entered in the competition by having performed certain intentional actions: filling out a form, or writing her name in a book, or whatever. Someone who wanders onto the competition grounds, and unintentionally releases an arrow, which strikes the target, has not in fact scored any points. For the same reason, someone who absent-mindedly moves chess pieces about, not knowing or intending to follow the rules of chess, even if her “moves” follow the pattern of a brilliant gambit, has not succeeded vis-à-vis the constitutive standards of the game of chess, simply because such a person is not playing chess in the first place.

The point here is nearly identical to that made in §2.1: the play a game, you have to *try to play*. But this kind of intentional trying is not plausibly a necessary condition on belief. Children and animals believe various things, but do not intend to engage in the “ordinary practice of

---

<sup>20</sup> Wedgwood 2002, p. 267-9. Wedgwood, however, is a normative essentialist (p.c. and 2007). The normativity of belief (e.g. the necessity of the truth norm) and the normativity of chess (e.g. the fact that such-and-such rules constitute the game of chess) are not analogous, on his view.

reasoning,” as Wedgewood would put it. And the same point about tacit<sub>1</sub> and tacit<sub>2</sub> intentions applies here as well.

### 2.3 Beliefs vs. artifacts

Essentialists sometimes analogize beliefs to artifacts, for example:

All sorts of things can “perform” well or ill when put to the test. Rational agents can do so, but so can ... designed instruments, and even structures with a function, such as a bridge. A bridge can perform well its function as part of a traffic artery. (Sosa 2009, p. 8)

I have been urging the following objection to performance essentialism: that believing does not always involve performance normativity, which it would need to do, to satisfy the Application Requirement, because believing is not paradigmatically based on a truth-directed intention (or any intention, for that matter). But perhaps this has been to misconstrue the requirements for “performance normativity” – the bridge has a constitutive standard of correctness, namely, safely supporting traffic as it passes over the river, but this depends not on the intentions of the *bridge!* The bridge need not try to support the traffic, for it to be a failure, qua bridge (or at least qua the particular bridge that it is), when it collapses.<sup>21</sup>

However, *someone* must intend that the bridge do so, for it to be a failure when it collapses. In most cases, the function of an artifact is determined by the intentions of the artificer – e.g. the engineers who design and build a bridge – but even in cases where something comes to have a function in virtue of being used a certain way – e.g. the old computer used as a doorstep – there must be someone who intends or wants the particular thing to function in a certain way. When a natural rock “bridge” collapses, it is no failure, and there is nothing bad about its collapse, in virtue of its failure, qua bridge. This is just to say that a bridge is an artifact, whose function supervenes on personal intentions (desires, goals, or whatever).

The problem with analogizing beliefs to artifacts is that beliefs are not analogous to artifacts in two ways, and both are ways that make the analogy unhelpful in explaining the legitimacy of epistemic evaluation.

First, the normativity involved in the evaluation of artifacts, in terms of their functions, is a very close cousin of performance normativity. If evaluation vis-à-vis the truth norm is the evaluation of beliefs in terms of their artifactual function, then someone must intend, of my belief that *p*, that it be true, or of my beliefs in general, that they be true, or of believers in general, that they have true beliefs. For naturalists who reject the idea of theistic design, none of these is plausible in the case of human beliefs. Unlike a bridge, intentionally created by engineers to support traffic as it passes over the river, my belief that *p*, and furthermore my beliefs in general, were not intentionally created by anyone. We rejected in §2.1 and §2.2 the idea that believers themselves always have truth-directed intentions, with respect to their beliefs. But there seems no one else – again, for naturalists who reject theism – in the neighborhood, as it were, upon whose intentions the artifactual function of belief might supervene.

---

<sup>21</sup> Thanks to Jonathan Ichikawa for suggesting this idea.

Second, there is a metaphysical problem with analogizing beliefs to artifacts. When engineers build a bridge, we are right to distinguish between their building of the bridge and the bridge itself. The bridge is a thing distinct from the act or event of its being built, and from the people who build it. In the sense of ‘belief’ relevant to the discussion of epistemic normativity, however, we would be quite wrong to think of beliefs as things distinct from believers, or as distinct from the act or event of their being believed. Some philosophers seem inclined to speak as though beliefs *were* analogous to bridges, in this respect, as though my having such-and-such beliefs were a matter of their being these separate things, my beliefs, to which I bear some relation of “having.” But this is wrong. When I believe that  $p$ , there is not this thing, my belief that  $p$ , to which I stand in some special relation. There is, perhaps, this thing, the proposition that  $p$ , to which I stand in the relation of believing. It is better to speak of believing, than of beliefs, and believing is something that I do (in a broad sense), and when I do it I stand in a certain relation to a certain proposition.

Consider, on the one hand, a “full blooded,” deliberate, and intentional action: having carefully considered the candidates for many weeks, I make up my mind, and on Election Day I cast my vote for so-and-so. Consider, on the other hand, an event that happens, and merely happens, inside my body: a pill that I have swallowed dissolves. Believing, it seems clear, belongs somewhere between these two extremes. I believe that T.E. Lawrence sometimes wore an Arabian robe. Up until right now, I’d never thought about the fact that I believed this. Having seen the movie, and having read a Wikipedia article about Lawrence, which featured a photograph of a man identified as Lawrence, wearing an Arabian robe, I came to believe that Lawrence sometimes wore an Arabian robe. I never wondered whether he really wore an Arabian robe, and I never considered the question of whether to believe this or not. What I’m saying seems clear is that, on the one hand, my coming to believe that Lawrence wore an Arabian robe, and my continuing to believe that he wore an Arabian robe, were neither of them “full blooded,” deliberate, intentional actions. I never had any intentions with respect of my believing something about Lawrence’s robes, nor did I ever deliberate about whether to believe that he wore Arabian robes. *If* we confine ourselves to my deliberate and intentional actions, then it was not something that I did, in that narrow sense. But, on the other hand, my believing that Lawrence wore Arabian robes was not something that just happened inside me, as it were, or something that happened to me. It is not as though my believing this is like a process that takes place inside me (like the dissolving of a drug), nor is my belief itself some kind of object that is inside me (e.g. in my head), the way that the dissolving pill is an object that is inside me.

Believing is, in this respect, like loving. Many philosophers have correctly seen that love is not a mental state, in the sense of being an object to which a person bears some intimate relation. Love is a relation, and loving is an activity; to love  $x$  is not to “possess” a mental state of love, with content  $x$ . It is to be disposed to act, emote, and think in the ways constitutive of loving  $x$ . Loving is an embodied activity (in a broad sense). But so is believing. (This does not mean, however, that believing is intentional, any more than it means that loving is.)

There is a perfectly legitimate sense in which beliefs are things whose existence is independent of the believer and her act of believing. This is sense of ‘belief’ on which this denotes the *content* of a person’s believing. ‘Katie and I have some of the same beliefs’ isn’t

the absurd claim that mine and Katie's minds overlap; it just means that there are some things believed by both me and Katie. 'We have the same belief' – meaning that we believe the same proposition.

#### *2.4 Conclusion*

Recall our dilemma (§2.1): the Application Requirement tells us that epistemic evaluation is legitimate in *all* cases of believing, so the conditions for epistemic evaluation had better be present in all cases of believing, but performance normativity requires intentions, or desires, or goals, on the part of the evaluated person, and it is implausible that believers always have the requisite intentions, desires, or goals. What we have seen suggests that there is no way for performance essentialism to evade our dilemma.

How then ought we respond to the sources problem? I have not considered the two other species of essentialism: biological essentialism and normative essentialism. (I also have not considered forms of epistemic essentialism other than those that embrace truth norm essentialism; but our dilemma would seem to threaten any account of epistemic normativity as performance normativity.) And I have not discussed the vices and virtues of the alternatives to epistemic essentialism, epistemic conventionalism and pragmatism about epistemic evaluation. But I believe that performance essentialism will not work, and we must look elsewhere for a solution to the sources problem.

## Bibliography:

- Alston, W.P. (1985), "Concepts of Epistemic Justification," *The Monist* 68, pp. 57-89
- Aristotle (2009), *Nicomachean Ethics*, trans. David Ross (Oxford University Press).
- Boghossian, P. (2003), "The Normativity of Content," *Philosophical Issues* 13, pp. 31-45
- (2005), "Is Meaning Normative?," in Nimtz and Beckermann (eds.), *Philosophy - Science - Scientific Philosophy* (Mentis).
- BonJour, L. (1985), *The Structure of Empirical Knowledge* (Harvard University Press).
- David, M. (2001), "Truth as the Epistemic Goal," in Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (Oxford University Press), pp. 151-69
- Davidson, D. (1982), reprinted in *Subjective, Intersubjective, Objective* (Oxford University Press, 2001), pp. 95-106.
- (1983), "A Coherence Theory of Truth and Knowledge," reprinted in *Subjective, Intersubjective, Objective* (Oxford University Press, 2001), pp. 137-53.
- Dennett, D. (1971), "Intentional Systems," *Journal of Philosophy* 68, pp. 87-106.
- Feldman, R. (1988), "Epistemic Obligations," *Philosophical Perspectives* 2, pp. 235-56
- (2000), "The Ethics of Belief," *Philosophy and Phenomenological Research* 60, 667-95
- Foley, R. (1987), *The Theory of Epistemic Rationality* (Harvard University Press).
- Goldman, A. (1986) *Epistemology and Cognition* (Harvard University Press).
- Greco, J. (2003), "Knowledge and Credit for True Belief," in DePaul and Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford University Press), pp. 111-34
- Grimm, S. (2008), "Epistemic Goals and Epistemic Values," *Philosophy and Phenomenological Research* 77, pp. 725-44.
- (2009), "Epistemic Normativity," in Millar, Haddock, and Pritchard (eds.), *Epistemic Value* (Oxford University Press), pp. 263-64
- Humberstone, L. (1992), "Direction of Fit," *Mind* 101, pp. 59-83
- Korsgaard, C. (1996), *The Sources of Normativity* (Cambridge University Press).
- Millikan, R. (1987), *Language, Thought, and Other Biological Categories* (MIT, 1987)

- Pritchard, D. (2007), "The Value of Knowledge," *Stanford Encyclopedia of Philosophy*, August 21, 2007.
- Railton, P. (1994), "Truth, Reason, and the Regulation of Belief," *Philosophical Issues* 5, pp. 71-93
- Riggs, W. (2002), "Reliability and the Value of Knowledge," *Philosophy and Phenomenological Research* 64, pp. 79-96.
- Shah, N. (2003), "How Truth Governs Belief," *Philosophical Review* 112, pp. 447-82
- Shah and Velleman (2005), "Doxastic Deliberation," *Philosophical Review* 114, pp. 497-534
- Sosa, E. (2001), "For the Love of Truth?," in Fairweather and Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility* (Oxford University Press), pp. 49-62
- (2007), *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1* (Oxford University Press).
- (2009), "Knowing Full Well: The Normativity of Beliefs as Performances," *Philosophical Studies* 142, pp. 5-15.
- Steglich-Petersen, A. (2006), "No Norm Needed: On the Aim of Belief," *Philosophical Quarterly* 56, pp. 499-516
- Velleman, J.D. (2000), *The Possibility of Practical Reason* (Oxford)
- Wedgwood, R. (2002), "The Aim of Belief," *Philosophical Perspectives* 16, pp. 265-97
- , R. (2007), *The Nature of Normativity* (Oxford University Press).
- Williams, B. (1976), "Deciding to Believe," in *Problems of the Self* (Cambridge), pp. 136-51
- Zagzebski, L. (1996), *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge University Press).
- (2003), "Intellectual Motivation and the Good of Truth," in DePaul and Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford University Press), pp. 135-54