

Epistemic Normativity

Stephen R. Grimm

How should we make sense of our epistemic evaluations? To judge a belief to be *justified* or *rational*, for example, is obviously to think something positive about it, and similarly to judge a belief to be *unjustified* or *irrational* is to think something negative. But what is the source or basis of these judgments?

Among contemporary epistemologists, perhaps the most prominent way to make sense of our epistemic evaluations is in teleological terms.¹ On this way of looking at things, a belief earns positive marks, from an epistemic point of view, just to the extent that it seems to promote or in some way bring about the things with intrinsic epistemic value. And similarly, a belief earns negative marks just to the extent that it seems to *fail* to promote or bring about the things with intrinsic epistemic value. I will say more about the motivation for this view in Section 1, but one of my basic goals in this paper will be to show that the teleological view—at least, as it is popularly understood—is mistaken.² In short, the problem for the view is that our practice of epistemic evaluation is broader and more wide-ranging than the view can capture. After considering a recent proposal by Ernest Sosa that seems to improve on the teleological account, I then suggest that Sosa's proposal too faces significant difficulties. I close by recommending a way of thinking about the nature of our epistemic evaluations that seems to avoid the problems canvassed earlier.

1. THE TELEOLOGICAL ACCOUNT

Although the teleological account is widely popular, in the first part of this paper I will focus on the way in which three philosophers in particular—Alvin

¹ In addition to the figures to be discussed below, see e.g. BonJour (1985: 7–8), Foley (1987: ch. 1), and Lehrer (1990: 112). For a more extensive list, see David (2001: 152).

² The qualification 'as it is popularly understood' will prove to be important later. There I will suggest that if one thinks about the goals that are being promoted (respected, etc.) in a different way, then it could be more plausible to think that we appraise our beliefs in terms of how well they promote these other goals and concerns.

Goldman, William Alston, and Michael Lynch—develop the view. Goldman, Alston, and Lynch are worth considering as a group, because although they eventually differ on the question of the value of true belief, they all begin at least by stressing the following two points. First, that as human beings we often value possessing certain epistemic goods *for their own sake*, and not merely for the sake of whatever further goals—especially, further practical goals—we might happen to have. And second, that the reason *why* we value these goods for their own sake is because of, or due to, our natural curiosity.

In the following passages Goldman, Alston, and Lynch not only endorse both claims, but help to show the way in which they seem to be naturally related:

[Goldman:] Our interest in information has two sources: *curiosity* and practical concerns. The dinosaur extinction fascinates us, although knowing its cause would have no material impact on our lives. We also seek knowledge for practical reasons, as when we solicit a physician's diagnosis or compare prices at automobile dealerships. (Goldman 1999: 3; emphasis added)³

[Alston:] [Although having true beliefs furthers our practical goals] the attainment of knowledge and understanding are *also of intrinsic value*. 'All men by nature desire to know,' said Aristotle, and this dictum has been reaffirmed by many of his successors. Members of our species seem to have a built-in drive to get to *the truth about things that pique their curiosity* and to understand how and why things are as they are and happen as they do. So it is as close to truistic as we can get in philosophy to take truth as a good-making characteristic, and falsity as a bad-making characteristic, of beliefs and other outputs of cognition. (Alston 2005: 31; emphasis added)⁴

[Lynch:] We care about the truth for more than just the benefits it brings us. . . . There are times in our lives when we simply want to know for no other reason than the knowing itself. *Curiosity* is not always motivated by practical concerns. Consider extremely abstract mathematical conjectures. With regard to at least some such conjectures, knowing their truth would get us no closer to anything else we want. (Lynch 2004: 15–16; emphasis added)

What I am calling the 'first' point therefore seems to be the more fundamental one: namely, that certain epistemic goods seem to possess a kind of intrinsic

³ This passage highlights the role of curiosity, but it is not as clear as it might be that Goldman associates this with a true-belief-for-its-own-sake claim. It is thus helpful to read this passage in conjunction with a passage from his earlier *Epistemology and Cognition* (1986). There he writes: 'Even if the desire for truth-acquisition is ultimately traceable to biological fitness (curiosity about one's environment can promote survival), it still appears in the organism as an 'autonomous' desire. People do not desire true belief merely as a means to survival, or the achievement of practical ends. Truth acquisition is often desired for its own sake, not for ulterior ends. It would hardly be surprising, then, that intellectual norms should incorporate true belief as an autonomous value, quite apart from its possible contribution to biological or practical ends' (Goldman 1986: 98).

⁴ Notice that, although Alston begins this passage by suggesting (along with Aristotle) that it is knowledge and understanding that is desired for its own sake, by the end of the passage (and in keeping with the rest of the argument in the book) he claims that it is truth that is the 'good-making characteristic'—in other words, the intrinsically valuable thing.

value. That is, they seem to be goods worth acquiring for their own sake, and not merely for the sake of whatever further practical goods they might help to produce.⁵ What I am calling the ‘second’ point, concerning the role of *curiosity*, in turn seems to be offered as a kind of explanation or defense of the first. After all, it might be thought, although little needs to be said on behalf of the idea that certain practical goods (such as pleasure, perhaps) are worth realizing for their own sake, it might be less obvious that epistemic goods have the same kind of status. Indeed, it might be thought that in comparison with other goods we obviously value for their own sake, the notion of a purely *epistemic* good seems like little more than a fiction.⁶ The appeal to curiosity, it seems clear, is meant to cut off just these concerns. Just as there are a range of practical goods we naturally desire, so too, the above passages suggest, there are also purely *epistemic* goods—goods the wanting of which can be explained in terms of our natural curiosity—that we naturally want to possess.⁷

But what are these ‘purely epistemic’ goods, exactly? In the passages quoted above there is not as much consistency as one might expect. Goldman first speaks vaguely of acquiring information and then of gaining knowledge, Alston first of acquiring knowledge or understanding and then more vaguely of something like possessing the truth, and Lynch first of caring for the truth and then of knowing the truth.

Despite this initial diversity, as we will see in a moment the considered view of all three seems to be that *believing the truth* is the thing that possesses intrinsic epistemic value, at least for creatures like us. When we are uncertain about how things stand with respect to certain subjects (Why *did* the dinosaurs die so suddenly, anyway?) our curiosity is naturally piqued by those subjects. Finding out the truth with respect to such subjects—in other words, believing the truth with respect to such subjects—accordingly possesses an intrinsic worth or value all its own.

⁵ Perhaps, as Goldman suggests, the nature of the dinosaur extinction is like this—though one would have thought the practical relevance of this topic (sudden and catastrophic extinction!) was fairly clear.

⁶ According to Stephen Stich (1990: 131), for example, although we do value many things intrinsically—health, happiness, the welfare of our children, etc.—the truth (Stich’s main candidate for a putatively epistemic good) is not one of them. Similarly, although Hilary Kornblith (2003: ch. 5) is critical of many aspects of Stich’s view, he seems to agree that there are no epistemic goods that are worth pursuing for their own sake.

⁷ In addition to the authors cited above, Ram Neta is another who makes this connection explicit. As he writes: ‘Knowledge and other positive epistemic statuses are worthy of pursuit by inquisitive creatures not (or not just) because they are instrumentally valuable. They may, of course, be instrumentally valuable—we need not disagree with Kornblith on that point. But that’s not the only thing that makes them worthy of pursuit for inquisitive creatures. What makes them worthy of pursuit for inquisitive creatures like ourselves is that, like health, friendship, and love, their attainment is partly constitutive of our well-being. Knowledge, and epistemic excellence more generally, is part of what constitutes the natural and valuable phenomenon of an inquisitive creature’s well-being’ (2007: 352).

Once we take believing the truth to be intrinsically valuable from an epistemic point of view, at any rate, for many philosophers the following teleological account of epistemic appraisal has come to seem very natural and compelling:

The teleological account of epistemic appraisal: A belief earns positive marks (counts as justified, rational, virtuous, etc.), from an epistemic point of view, just in case it does well with respect to the things with intrinsic epistemic value (i.e. helps to promote them or bring them about). Likewise, a belief earns negative marks just in case it does poorly with respect to the things with intrinsic epistemic value.⁸

As the following passages suggest, Goldman, Alston, and Lynch all endorse this view in very similar terms, and make it clear (or, at least, clearer) that by their lights true belief is the thing with intrinsic epistemic value.⁹

[Goldman:] I shall attempt to make a case for the unity of epistemic virtues in which the cardinal value, or underlying motif, is something like true, or accurate, belief. . . . The principal relation that epistemic virtues bear to the core epistemic value will be a teleological or consequentialist one. A process, trait, or action is an epistemic virtue to the extent that it tends to produce, generate, or promote (roughly) true belief. (2002: 52)

[Alston:] We evaluate something epistemically (I will be mostly concerned with the evaluation of beliefs) when we judge it to be more or less good from the epistemic point of view, that is, for the attainment of epistemic purposes. . . . The evaluative aspect of epistemology involves an attempt to identify ways in which the conduct and the products of our cognitive activities can be better or worse vis-à-vis the goals of cognition. And what are those goals? Along with many other epistemologists I suggest that the primary function of cognition in human life is to acquire true beliefs rather than false beliefs about matters that are of interest to us. (2005: 28)¹⁰

⁸ This account is doubtless incomplete as it stands (though complete enough for our purposes), because presumably *withholdings* too can be appraised epistemically. For more on this see DePaul (2004).

⁹ I use the notion of 'intrinsic value' in what I take to be the standard way here, to mean a value that is worth pursuing and realizing for its own sake. When Goldman speaks of truth as a 'cardinal value' we pursue for its own sake, I therefore assume by this he means what we standardly mean by an 'intrinsic value.' (It is possible that by the talk of 'cardinal' instead of 'intrinsic' value, however, Goldman has something more like the Sosa view, which we will consider shortly, in mind.) Similarly, Lynch (elsewhere) prefers to speak of the for-its-own-sake value that believing the truth possesses as a 'constitutive value' (see e.g. Lynch 2004: 127)—'constitutive' in the sense that it is an essential constitutive part of a flourishing life, which is an end we all desire. As Lynch notes, the notion of a constitutive value is theoretically quite similar to the notion of an intrinsic value: 'Being constitutively good, like being an intrinsic good, makes something worth caring about for its own sake, as opposed to caring about it for what it leads to' (2004: 128).

¹⁰ Notice that, unlike Goldman (and Lynch, looking ahead), Alston claims not simply that true belief is the intrinsically valuable thing, but 'true beliefs rather than false beliefs about matters that are of interest to us.' This is not a trivial difference, as I will argue at length in Section 3, but for the moment we can put the distinction to one side.

[Lynch:] Once again, the key point is that the value of believing what is justified is parasitic on the value of believing what is true. Having justified beliefs is good because justified beliefs are likely to be true. (2004: 50)

Although more would need to be said here about many aspects of the view,¹¹ I hope that by now the basic idea is clear enough: again, that the reason why we think of an individual belief as good or bad is because of some sort of ‘doing well’ relationship that the belief bears towards the things with intrinsic epistemic value, where the things with intrinsic epistemic value are taken to be true beliefs.

2. TWO VIEWS OF VALUE

We just saw that Goldman, Alston, and Lynch all motivate their views by pointing out that, when we are uncertain about how things stand with respect to certain subjects, our curiosity is naturally piqued by those subjects. Suppose for the moment we grant that some subjects *do* in fact naturally elicit our curiosity: perhaps something like the dinosaur extinction falls into this category. Finding out the truth with respect to these subjects will then seem to be intrinsically worthwhile, from a purely epistemic point of view.

The question we now need to ask, however, and the one that reveals an instability in the teleological view, is whether it is really plausible to think that just *any* subject falls into this category. Suppose I am uncertain about how many motes of dust there are on my desk now, for example, or about the now-defunct phone number of some random person in Bangladesh. Is my curiosity really naturally elicited by *these* subjects? Are these really the kinds of subjects that Aristotle had in mind when (as Alston notes in his earlier passage) he claimed at the outset of the *Metaphysics* that ‘All men by nature desire to know’?

Significantly, this is where opinions begin to divide. According to Lynch,¹² for example, finding out the truth with respect to just *any* subject—even apparently trivial subjects like the number of motes of dust on my desk—possesses genuine, intrinsic epistemic value: it possesses a value worth pursuing for its own sake, from a purely epistemic point of view. Of course, Lynch is quick to acknowledge that the value to be found in trivial subjects of this sort is usually trumped by our other concerns—the value is therefore only *prima facie*, by his lights. But on his view, and had we world enough and time, finding out the truth with respect to any of these topics would indeed be intrinsically worthwhile, from a purely epistemic point of view.¹³

¹¹ For example, what sort of reliability matters? ‘Actual world’ likelihood (where ‘actual’ is a name, rather than an indexical)? ‘Normal world’ likelihood? Something else? This is none too clear, as Goldman’s various stances over the years suggest.

¹² Among others: Kvanvig (2003: 41) and Horwich (2006: 347) also defend this view.

¹³ As Lynch, responding to the sort of natural objection we will next consider, writes: ‘Come on, what about really trivial truths? Surely there are all sorts of true beliefs I could have that are not

Let's call this the *unrestricted* view of the value of true belief, according to which believing the truth with respect to just *any* subject possesses a kind of intrinsic epistemic value.¹⁴ The reason why opinions begin to divide here is that many philosophers have found the unrestricted view of the value of true belief incredible, including (most notably) fellow advocates of the teleological account such as Goldman and Alston.

Goldman offers his own counterexamples that tell against the unrestricted view (1999: 88; 2002: 61). What is the 323rd entry in the Wichita, Kansas, phone directory? Who was placed sixth in the women's breast stroke in the 1976 Summer Olympics? What was the full name of Domenico Scarlatti's maternal grandmother? According to Goldman, since finding out the answer to questions of this sort seems wholly lacking in value—even, it seems, from a 'purely' epistemic point of view—'We can no longer suggest that higher degrees of truth possession are all that count in matters of inquiry' (2002: 61). Instead, on his view, we need to shift to what we might now call a *restricted* or *qualified* view of intrinsic epistemic value, where what matters is not possessing the truth on any topic but rather only on 'topics of interest' (2002: 61; cf. Goldman 1999: 89).

In his (2005) book Alston too argues that an unrestricted view of the value of true belief cannot be maintained.¹⁵ Since the true beliefs that we could gain from activities like memorizing phone books apparently lack intrinsic value, Alston suggests, along with Goldman he concludes that we need to restrict the realm of those things with intrinsic epistemic value to truths concerning 'matters that are of interest or importance to us' (2005: 32).¹⁶

Despite their initial agreement, the advocates of the teleological account we have been considering so far therefore part ways conspicuously when it comes to identifying the thing (or things) with intrinsic epistemic value. Although all begin with the claim that true beliefs are the things with intrinsic epistemic value, in the face of certain obvious objections—especially, what we might think of as the 'trivial truths' objection—Goldman and Alston immediately back off their claim and relativize the intrinsically valuable things to true beliefs on, roughly, 'matters of interest or importance to us.'

even *prima facie* good? Without a doubt, there are all sorts of true beliefs that are not worth having, all things considered. But the fact that I should not bother with those sorts of beliefs doesn't mean that it isn't still *prima facie* good to believe even the most trivial truth' (2004: 55).

¹⁴ Lynch puts his point more formally as follows: 'It is *prima facie* good, for all p (to believe that p if and only if it is true that p)' (2005: 331).

¹⁵ It is worth noting that his 2005 book represents something of a change in his thinking about the nature of epistemic value. In his earlier 'Concepts of Epistemic Justification,' for example, he there characterized the epistemic goal as that of 'maximizing truth and minimizing falsity in a large body of beliefs,' although he explicitly called that a 'rough characterization' (1989a: 83–4).

¹⁶ And, indeed, many epistemologists writing on the topic of our 'epistemic goal' quite naturally relativize the goal to something like 'topics of interest and importance.' See e.g. Haack (1993: 199) and David (2005: 299). The temptation to make this move is obviously very powerful.

But how dramatic is this difference, and what does it have to teach us about the viability of the teleological view? If Goldman is to be believed, the qualification represents only a 'slight' revision to the teleological view. As he writes:

But can't we incorporate the element of interest by a slight revision in our theory? Let us just say that the core epistemic value is a high degree of truth possession *on topics of interest*. Admittedly, this makes the core underlying value a somewhat 'compound' or 'complex' state of affairs. But, arguably, this is enough to preserve the idea of thematic unity, and thereby preserve Unitarianism. (2002: 61)¹⁷

Alston too seems to think that the revision is quite slight; at any rate, he seems to even lack Goldman's misgiving that such a qualification immediately 'makes the core underlying value a somewhat "compound" or "complex" state of affairs.'¹⁸

What I want to argue in the following section, however, is that this difference concerning the nature of the intrinsic epistemic value is dramatic indeed, and that it exposes a fundamental problem at the heart of the teleological view.

3. A DILEMMA

To see why, suppose we take it, along with Goldman and Alston, that not all true beliefs are intrinsically valuable but only true beliefs with respect to subjects of interest or importance to us. This then leads us to the crucial question: How should we make sense of our epistemic appraisals with respect to those beliefs (or, better, those topics) that apparently *lack* intrinsic epistemic value—that is, that are not interesting or important, from a purely epistemic point of view?

If we take the teleological account at its word, such a belief would deserve a positive or negative appraisal only to the extent that it did well with respect to the things with intrinsic epistemic value. But by hypothesis a true belief on such a topic would *lack* any such value.¹⁹ And from this it would seem to follow that a positive or negative appraisal of the belief would simply be out of place.

But now the problem should be clear, for positive and negative appraisals of such beliefs clearly do *not* seem out of place. Suppose that on a lazy whim you decide to scan your desktop for motes of dust. After a bit of distracted counting, you then conclude that the desktop is harboring eighteen motes. Given the sloppiness of your method, however, we can suppose that this answer really amounted to little more than a guess: you might very easily have concluded, for example, that there were rather more motes or rather fewer.

¹⁷ 'Unitarianism' is Goldman's term for the view that the only intrinsic epistemic value—in his words, the 'cardinal' epistemic value—is true belief.

¹⁸ See e.g. Alston's discussion in 2005: 30–3.

¹⁹ Some philosophers even insist (on a slightly different note) that a *false* belief on this topic holds no intrinsic disvalue. See e.g. Kelly (2003: 624–5).

What now should we say about your belief? Even if it turns out to be true, is it justified? Well-formed? Rationally held? I take it that on all counts the answer is No. Given your lack of responsiveness to the truth about the moles, your belief would presumably earn low marks with respect to virtually any type of epistemic appraisal on offer.

The problem for the restricted teleological account²⁰ offered by Goldman and Alston, however, is to explain why this should be so. Recall that according to the teleological account a belief inherits its epistemic status from its ‘doing well’ relationship to the things with intrinsic epistemic value. But if there is nothing with intrinsic epistemic value to do well with respect to, then it is hard to see where this inherited value or status might come from. From nothing, nothing comes, it would seem.

This then leads us, by way of summary, to the following dilemma for the teleological account of epistemic appraisal. For suppose that, with Goldman and Alston, it is *not* the case that believing the truth with respect to just any subject is intrinsically worthwhile, from an epistemic point of view. If so, then the teleological account seems unable to explain why our beliefs with respect to subjects that lack this kind of value—for short, ‘trivial’ beliefs—are appropriate candidates for epistemic appraisal. Suppose instead that along with Lynch, we accept an unrestricted view of intrinsic epistemic value. Combined with the teleological account, we can then make sense of the fact that ‘trivial’ as well as ‘non-trivial’ beliefs are appropriate candidates for epistemic appraisal, for such beliefs would be derivatively either good or bad to the extent that they promote or respect the things with intrinsic value. But, again, the problem with this account is that an unrestricted view of intrinsic epistemic value is deeply implausible. It hardly seems to be the case that finding out how things stand with respect to just any subject is intrinsically worthwhile, even from a purely epistemic point of view.²¹

In the following section I will consider a recent proposal by Ernest Sosa that seems to allow us to keep the spirit of the teleological account while avoiding the sorts of problems that arise when we try to take a stand on a particular account of intrinsic epistemic value. Before moving on, however, I want to consider one way that Goldman and Alston might try to blunt their particular horn of the dilemma just described.

Recall that Goldman and Alston both qualified their initial views by restricting the scope of intrinsic epistemic value to (roughly) ‘matters of interest or importance.’ But now it might be thought that the sort of problem cases imagined above—as when we negatively appraise my sloppy counting of moles

²⁰ That is, the problem arises for an account that combines the teleological view of appraisal with a restricted view of intrinsic epistemic value.

²¹ For more on the selectiveness of our sense of curiosity, see Harman (1999: 100) and Grimm (2008).

of dust—are not *really* problems because they are not fairly described. As I described the mote-counting case, for example, I gave the impression that I could care less about how things stood with respect to this subject; instead, I was just looking to pass the time. But it is implausible to suppose that anything like a genuine *belief* could issue from such a process, for the very process of forming a belief seems to require that I care about how things stand with respect to the subject in question.²² Thus when a topic *interests* me enough to trigger a belief—in other words, when I care about it enough—it might be thought that this implies the presence of something *worth* caring about, hence the presence of a value that could be used to ground further epistemic appraisals.

The basic problem with this response, however, is that it loses track of the fact that not just any sort of good was supposed to ground the teleological account, but rather a good that was distinctively ‘epistemic’; in other words, a good that we took to be intrinsically worth realizing from a ‘purely’ epistemic point of view, or (apparently) simply insofar as we were *curious* beings. Assuming the sort of interest described in the desktop case was serious enough to issue in a belief, however, it was nonetheless still at bottom a practical interest—it was an interest that stemmed from my desire to put off my work for a little while longer, or to give my mind a little rest, and so on. It was, presumably, not an interest that derived from a purely (or even partly) curiosity-driven inclination of mine.

The teleological view, at least as developed by philosophers such as Goldman and Alston, therefore has significant problems trying to make sense of how our concepts of epistemic appraisal apply to apparently ‘trivial’ topics. But perhaps the teleological view is better understood in a different way—one that retains the appealing structural features of the view while bypassing the difficulties that surround the question of epistemic value. As we will see in the following section, Sosa’s recent work offers just such an alternative.

4. SOSA’S VIEW

Sosa begins his account by noting that as human beings we are ‘zestfully judgmental’ across a wide range of areas, including art, literature, science, politics, sports, food, wine, and even coffee (2007: 70). In Sosa’s terms, each of these areas of evaluation therefore represents a kind of *domain*—more exactly, each area represents a *critical* domain. Why ‘critical’? Because, Sosa suggests, once we identify the values that are *fundamental* within each domain we can then appraise or assess (hence criticize) the *derivative* value of other items in the

²² See Nishi Shah (2003), and Shah and David Velleman (2005). It is also true that a typical person will have what we might think of as ‘standing concerns’: thus if I hear a loud noise nearby, or see a flash in the distance, I will (as it were) automatically form a belief about these subjects. Thanks to Robert Audi for helping to clarify this point.

domain in terms of how well they promote, bring about, or perhaps in some way duly respect the domain's fundamental values. The fundamental values within a given domain therefore serve as the goal around which the critical domain is structured.

Consider, for example, the domain of assassinship. Because for an assassin killing one's target is the goal around which the practice of assassinship is structured (or so it seems), we can therefore evaluate various elements of the assassin's conduct in terms of how effectively the conduct realizes this goal. Or consider the critical domain associated with the card game whist. Since the goal of whist is to take the majority of tricks, particular moves in whist can therefore be evaluated in terms of how well they realize *this* fundamental goal.

Moreover, what focusing on these more unusual sorts of domains helps to bring out, Sosa suggests,²³ is that our ability to evaluate particular items within a domain does not turn on our judgments about the worth of the fundamental values that structure the domain. Thus with respect to some domains, such as whist playing, we might think that the fundamental values involved are too trivial to possess any intrinsic worth. And with respect to others, such as assassinship, we might even think that the fundamental values that structure the domain possess positive *dis*value. Nonetheless, as Sosa notes, this hardly seems to get in the way of our ability to appraise particular elements within the domain:

Paradoxically, one can be an adept critic within such a domain even while discerning no domain-transcendent value in it. Thus, someone knowledgeable about guns and their use for hunting, for military ends, etc., may undergo a conversion that makes the use of guns abhorrent. The good shot is thus drained of any real value that he can discern. Nevertheless, his critical judgment within that domain may outstrip anyone else's, whether gun lover or not. Critical domains can be viewed as thus *insulated*, in ways suggested by our example. (2007: 73–4)

The basic insight Sosa wants to build on, then, is that we can evaluate items within a domain in terms of how effectively they promote or bring about the fundamental values of the domain, while all the while remaining agnostic about whether the domain's fundamental values are valuable or worth pursuing intrinsically.

So, how does this basic insight help to shed light on our concepts of epistemic appraisal? According to Sosa, epistemic appraisals too take place within an insulated critical domain, a domain in which the fundamental value is true belief. Unlike Goldman, Alston, and Lynch, however, Sosa argues that in order to make sense of our concepts of epistemic appraisal there is no need to take a stand on whether true belief is something that possesses intrinsic value or is worth pursuing for its own sake. In Sosa's words: 'Truth may or may not be intrinsically

²³ I should add that the examples just mentioned are meant to *illustrate* Sosa's view; they are not Sosa's own.

valuable absolutely, who knows? Our worry requires only that we consider truth the epistemically fundamental value, the ultimate explainer of other distinctively epistemic values' (2007: 72).

5. A CLOSER LOOK

What should we make of Sosa's proposal? On the positive side, the view naturally accommodates our ability to appraise beliefs on *any* topic, even apparently 'trivial' topics. Since the fundamental epistemic value for Sosa is simply true belief, it follows that 'trivial' beliefs can be appraised and evaluated just as readily as more 'important' beliefs. In this way the account captures the full scope of our epistemic appraisals; unlike the Goldman–Alston view, it doesn't leave the trivial out. The view can also make good sense of our appraisals concerning how *effectively* a particular believer reaches the truth goal. Although the terms he uses are a bit Sosa-specific, there clearly seems to be a sense in which we can evaluate a belief as 'adroit' or as 'maladroit' (i.e. as deriving, or not, from a reliable competence on the part of the believer to realize the truth), or as 'apt' or 'inapt' (i.e. as realizing, or not realizing, the truth *because of* such a competence).²⁴ Sosa's view therefore interestingly unites our appraisal of beliefs to our appraisal of performances more generally. Thus, just as we can judge an archer's shot to be adroit or maladroit (relative to the goal of striking the bullseye) or we can judge a tennis player's serve as apt or inapt (relative to the goal of hitting the ball in the appropriate box), so too we can judge the truth-oriented merits of someone's believing: as a performance that manifests various degrees of skill and efficiency relative to the truth goal.

Despite these virtues, what I want to suggest now is that by remaining agnostic about the domain-transcendent value of true belief, Sosa seems to introduce a new problem—seems to, indeed, lose sight of one of the most important aspects of our epistemic appraisals. For notice: when we judge a belief to be unjustified or irrational, we seem to be doing more than just evaluating (in this case, in a negative way) the skill or virtuosity of the believer's performance. In addition, we seem to be in some sense *criticizing*, perhaps even *reproaching*, them for believing in this way.²⁵ To judge someone's belief to be unjustified or irrational is thus to judge that the person's attitude towards the content of the belief *should* be

²⁴ Sosa has been working with this distinction for some time; see e.g. Sosa 1991. For his most recent version, see his 2007: ch. 2.

²⁵ As Nicholas Wolterstorff notes: 'We say to each other such things as, "You should have known better than to think that Borges was an English writer," "You should be more trusting of what our State Department says," and "You should never have believed him when he told you that the auditors had approved that way of keeping books." Not only do we *regret* the knowledge and ignorance of other human beings, their beliefs, disbeliefs, and non-beliefs; we reproach them, blame them, chastise them, using the deontological concepts of ought and ought not, should and should not. Of course we also praise them for believing and not believing, knowing and not knowing, as they do' (2005: 326).

reconsidered, in some apparently binding sense of ‘should.’²⁶ As Hilary Kornblith puts the point:

If you tell me that a belief of mine is unjustified, this gives me reason to give up that belief. The epistemic claim is something about which I should care, and an account of the source of epistemic norms must explain why it is that I should care about such things. (2002: 145)

What’s more, even what we referred to earlier as the Sosa-specific appraisals such as ‘apt’ and ‘inapt’ seem to carry with them this normative force. Thus, and to extend Kornblith’s point, if I accept that a certain belief of mine is ‘inapt’ I seem now to have a reason to do something about my attitude toward the content of the belief: perhaps to change my attitude altogether, or perhaps to try to get into a better epistemic position with respect to the subject at issue, and so on. Simply sticking with the original attitude in the face of the ‘inapt’ judgment does not seem to be acceptable. If at all possible, it seems that I *should* try to do something about my position.²⁷

Can Sosa make sense of the way in which this binding sense of ‘should’ attaches to our epistemic evaluations (especially, it seems, our negative evaluations)? In one sense it might be thought that he can, for there does seem to be a natural place for a ‘should’ even within Sosa’s insulated domains of critical appraisal. Thus, we might say that, given that such-and-such is the goal, one *should* proceed in this way—and not in that way—in order to realize the goal. So, for example, given the goal of acquiring a true belief with respect to a given subject, one *should* base one’s belief on good evidence, rather than hazard a random guess, because basing one’s belief on good evidence is a more effective way of realizing the goal at issue.

²⁶ Notice that in suggesting that beliefs are subject to criticism in this way we do not have to accept that belief is subject to our direct voluntary control. I think (along with Alston 1989*b*, Plantinga 1993, and virtually everyone else) it is obvious that it is not. Instead it seems that all we need suppose is that belief is under *enough* control to make critical judgments appropriate. At any rate, these types of judgments are central enough to our epistemic appraisals that any theory of epistemic normativity should seek to accommodate them (see e.g. Audi 2001, as well as the previous footnote). Notice as well that while I have argued that a judgment that a belief is (say) truth-unjustified carries with it the judgment that the subject of the belief should try to improve her cognitive position with respect to the belief, I do not mean to say that it is always psychologically possible to give up the belief. Sometimes, as a result of brainwashing, perhaps, or possible psychological trauma, it might not be. But it does not follow that the ‘unjustified’ judgment does not have this binding sense of ‘should’ attached to it. Consider a comparison with the judgment *cruel*. I take it that when we judge a particular action to be cruel, this brings with it the idea that the agent should not act in this way, even that the agent has a binding reason not to act in this way. But it seems equally clear that we might apply this judgment to the actions of a particular agent even if, for some peculiar psychological reasons, the agent felt compelled to act this way, perhaps to the point where he could not have acted otherwise. Indeed, in a loose, analogical way, we sometimes even apply the judgment ‘cruel’ to the behavior of animals, even though it seems unlikely that they have the sort of voluntary control over their actions that would make judgments of blame and censure strictly appropriate.

²⁷ Alan Millar (2004: 92–9) and Terence Cuneo (2007: 67–70) likewise tie the notion of normativity to the notion of *having a reason*.

The problem of course is that this sense of ‘should’ is quite weak; in roughly Kantian terms, it is the ‘should’ of calculation rather than the stronger, binding ‘should’ associated with duty or obligation. In this weaker sense of ‘should,’ after all, the assassin *should* use a high-powered rifle, rather than a flimsy slingshot, in order best to realize his goal of killing the target. The sense of ‘should’ associated with our judgment that a particular belief is unjustified, however, seems deontologically more substantial than that.

To make better sense of the binding sense of ‘should’ that seems to attach to our epistemic appraisals, it is worth recalling that we can appraise or evaluate beliefs relative to several different goals.²⁸ Suppose, for example, I am wondering whether God exists. It might be the case that if I were to believe that God did *not* exist then I would experience tremendous psychological distress: I would find it very hard to go on in a world that suddenly seemed devoid of meaning. We can therefore appraise how well my belief about God does not just with respect to the goal of realizing the truth but also with respect to this other goal—roughly, what we might think of as the goal of ‘psychological comfort.’

Imagine now that after soberly weighing the evidence I decide that God does not exist, thus (as expected) bringing with it significant psychological distress. Relative to the goal of psychological comfort (‘comfort,’ for short), we can therefore appraise my belief in a variety of different ways: we can say that it was ‘comfort unjustified,’ ‘comfort irrational,’ ‘comfort inapt,’ ‘comfort maladroit,’ and so on. Suppose we settle on one of these judgments: that the belief was ‘comfort unjustified’—that is, unjustified from the point of view of psychological comfort. If I accept this judgment, does it now follow that I have *reason* to give up my belief, or that I *should* give up my belief? I take it that in some weak, calculative sense of a ‘reason’ or of ‘should’ this might be right. Thus, relative to the goal of psychological comfort, I have a reason to give up my belief; alternatively: relative to this end, I should give it up. But it seems clear that neither the ‘reason’ nor the ‘should’ at issue here is binding in the way considered above—in neither case does it seem that we can be justly blamed or criticized for failing to orient our belief towards the goal of psychological comfort, for example.

Once we consider things from the truth perspective, however, we can see the normative force of our evaluations has a dramatically different character. Suppose, for example, that instead of soberly weighing up the evidence I formed my belief about God by hazarding a random guess. Relative to the goal of believing the truth, naturally, this belief will earn a variety of negative appraisals: thus we might say (extending our artificial evaluations for the moment) that the belief was ‘truth unjustified,’ ‘truth irrational,’ ‘truth inapt,’ ‘truth maladroit,’ and so on. Notice now, however, that if I were to agree that my belief was

²⁸ I do not want to suggest that this point is unfamiliar to Sosa (indeed, he begins many of his papers with the very distinction)—just that its force needs to be appreciated properly here.

truth-unjustified, for example, I would now have more than a calculative reason to give up my belief, a reason that would be potentially dispensable, if for some reason I no longer cared about the truth. Instead, I would seem to have a *binding* reason. In accepting this judgment, I would agree that I should not be holding this belief, in some non-optional sense of ‘should.’

It seems clear enough, then, that even though we can evaluate beliefs relative to countless different ‘fundamental values’ and hence countless different domains, the end of realizing the truth enjoys a special sort of status when it comes to the evaluation of belief. What’s more, the fact that realizing the truth enjoys this special status seems to account for the particular normative force that our epistemic appraisals possess. Given Sosa’s agnosticism about the domain-transcendent value of true belief, however, it is not clear that he can make sense of the fact that the truth perspective is in some sense the privileged perspective—as we might say, the binding perspective—when it comes to the evaluation of belief.

Before moving on, I should note that, towards the end of his most recent discussion of epistemic normativity, Sosa offers a distinction that might seem to accommodate the sort of bindingness we’ve just been emphasizing. Thus Sosa suggests that we need to distinguish between two sorts of epistemic normativity: on the one hand, the sort of normativity that is *constitutive* of knowledge, and on the other hand, the sort of normativity that is relevant to the study of ‘intellectual ethics’ (2007: 89–91). On this way of looking at things, the normativity associated with ‘intellectual ethics’ has to do with appreciating and pursuing the sorts of intellectual topics that are in some sense ‘finer’ (2007: 89) and hence more *worthy* of our attention. But according to Sosa this sort of normativity—which would seem to bring with it a kind of binding character—has little if anything to do with the sort of normativity that is *constitutive* of knowledge. Thus while we might blame or criticize someone for spending their life counting blades of grass (say), these judgments are separable from our normative judgments about whether a particular belief counts as apt or inapt, adroit or maladroit, and so on.

Although Sosa seems right that there is an important distinction to be made here, the question we still need to ask is whether it helps to explain the distinctive normative force of our epistemic appraisals. And it seems to me that it does not. As I noted at the outset, in trying to offer an account of epistemic normativity we are presumably looking for an explanation of why our epistemic appraisals have the particular force that they evidently have. I have argued, moreover, that our epistemic appraisals have a particular binding or reason-giving force that other sorts of appraisals lack (say, appraisals concerning comfort-aptness or inaptness). But Sosa’s distinction seems to offer no explanation of this last fact. Even if we grant him (as I think we should) that a judgment that S should not be counting blades of grass seems categorical rather than merely hypothetical, we still have no explanation for why a judgment that S’s grass belief was maladroit (say) goes hand in hand with a judgment that S should try to improve his cognitive position

with respect to the grass, if possible—where the ‘should’ here again seems to be a categorical one, rather than a merely hypothetical one.²⁹ It seems that we need to look further, then, in search of the source of the special sort of normativity that attaches to our epistemic appraisals.

6. OUR POSITION

Let’s take stock. Recall that Section 3 closed with a dilemma for the teleological account of epistemic appraisal. If we suppose that only truths of interest or importance are intrinsically valuable (from a ‘purely’ epistemic point of view—whatever exactly that comes to), then it looks like we lose our ability to explain how our epistemic appraisals apply to unimportant, ‘trivial’ beliefs. If we suppose instead that *any* truth is intrinsically valuable, then it looks like we’ve reached an absurdity; only someone really desperate, apparently, would think *that*. Our discussion of Sosa’s view, in turn, suggested that attempts to make sense of our epistemic appraisals should not lose sight of the distinctive normative force of these appraisals. Thus to judge a belief to be justified (for example) is not simply to judge that it is skillfully oriented to the truth but rather that it *should* be so oriented, in some binding sense of ‘should’—just as to deem a belief to be unjustified is to judge that it should *not* be so oriented, in some binding sense of ‘should not.’ To suppose that it is only, or even primarily, the skillfulness of the belief that we are appraising when we make positive appraisals of this kind seems to lose sight of the fact that a belief can be skillfully aimed at almost any goal. There thus seems to be something special about the truth goal that Sosa’s truth agnosticism apparently leaves out.

Overall, this leaves us with two main points in need of reconciliation. First, the fact that our epistemic appraisals not only can be, but patently are, applied to the full range of our beliefs. Second, that our epistemic appraisals seem to have a distinctive normative (binding) force, suggesting that the truth goal is not simply one goal among others when it comes to the evaluation of belief, but rather a goal with a special status.

What I want to propose in this section is that the best way to make sense of these two claims is to accept a *modified* version of the thesis that any true belief has a special value or worth—a version that attempts to explain the unrestricted value of true belief in a slightly different way, or (better) from a slightly different perspective. More exactly, I want to argue that the best way to make sense of the two claims is by shifting away from the standard first-person question about the value of true belief—wherein we ask (as Goldman, Alston, and Lynch

²⁹ It is worth noting that Sosa’s position, as far as I can see, does not exclude the possibility of identifying a further source for the bindingness that characterizes our epistemic appraisals. The problem is only that the view as it stands leaves this further sort of normativity unexplained.

asked) about the value of true belief in terms of our own intellectual goals or well-being—and by moving instead towards a more communal or social view of the value of truth.³⁰

To appreciate how the switch from a first-person perspective to a social perspective might help, consider first the following passage from Thomas Kelly. Here Kelly makes a point that will by now be familiar: namely, that it hardly seems to be the case that believing the truth with respect to just any subject is worthwhile or valuable, even from a ‘purely epistemic’ point of view. More radically, Kelly also insists that even believing something *false* about many subjects does not bring with it any obvious disvalue. As he writes:

In addition to those many truths such that my believing them would contribute to the achievement of some goal that I have, there are also (countless) truths such that my believing them would not contribute to any goal that I actually have. Whether Bertrand Russell was right- or left-handed, whether Hubert Humphrey was an only child—these are matters of complete indifference to me. That is, I have no preference for having true beliefs to having no beliefs about these subjects; nor, for that matter, do I have any preference for having true beliefs to false beliefs. There is simply no goal—cognitive or otherwise—which I actually have, which would be better achieved in virtue of my believing true propositions about such subjects, or which would be worse achieved in virtue of my believing false propositions about them. (2003: 624–5)

Let’s grant for the moment that Kelly is right about this: that believing the truth with respect to such topics would not contribute to any personal goals he might happen to have. What’s it to him, then, if he forms a belief about whether (say) Humphrey was an only child by means of a random guess? And yet, as we saw before, such a random guess would not only earn low marks from an epistemic point of view (count as unjustified, irrational, etc.), but also earn our criticism and perhaps even reproach. Why so?

Considered from a social point of view, it would seem that the answer to this question can be found by noting that even though we might not care less about some belief (or better, some topic), it is nonetheless the case that other people might care about the topic a great deal. For example, while finding out the truth with respect to whether Humphrey was an only child may not hold any value for us, or may not elicit our curiosity in any way, presumably for Humphrey’s biographer (say) getting this right will be quite important—if not in itself, or for its own sake, then at least for the sake of producing an accurate account of Humphrey’s life. And given that someone in the biographer’s shoes might depend on us as potential sources of information about this topic, it seems

³⁰ It might sound odd to accuse Goldman, at least, of being insensitive to the social value of true belief, for perhaps more than any current epistemologist he has emphasized the importance of the social dimension of knowing! To my mind, however, his teleological framework of epistemic appraisal does not sufficiently reflect this fact. Thus, for example, his basic framework is essentially indistinguishable from Alston’s, who does *not* stress the social in the same way.

that we have an obligation not to be cavalier when we form beliefs about the question—in other words, an obligation to try to position ourselves well with respect to this question, even though, given our own cognitive goals, we might very well be thoroughly uninterested.³¹

In short, what turning away from one's personal goals and concerns and towards our broader role in the information economy (as it were) helps to remind us is that the concerns of others—especially the practical concerns of others—are remarkably plastic and unpredictable. As such, it reminds us that, even though a topic may hold no interest or value from our first-person point of view, it may well hold interest or value for someone else. And as a potential source of information for others, we have an obligation to treat *any* topic or any question with due respect.³²

Suppose this approach works with respect to the Humphrey question. But what about the *really* trivial topics, like the 323rd entry in the Wichita, Kansas, phone directory? Or the number of motes of dust on my desk at this very moment? Although it is harder to imagine how someone might have an interest in these questions, once again we need only bear in mind how odd and varied people's practical concerns can be. For the Wichita phonebook fact-checker, for example, it might well be important—not epistemically important, but practically, presumably—to know the name of the 323rd person in the directory. Less realistically, but still possibly, we can imagine that someone with a particular antipathy towards dust might well wonder whether his new 'anti-dust' strategies have really succeeded in cutting down the number of motes of dust on his desk, as he fondly hopes.

It therefore seems that the basic idea we need to make sense both (a) of the apparently unlimited range of our epistemic appraisals as well as (b) of their normative force is that, given our nature as information-dependent and information-sharing creatures, we have an obligation not just to be sources of information for others but to be *good* sources of information. This obligation stems, moreover, not from the fact that believing the truth with respect to just any subject is intrinsically valuable, but rather because any subject might *come* to have value—if only value of a practical sort—in light of the varied and unpredictable concerns others might have.³³

Plausibly, then, the best way to make sense of the value of true belief is to think of it along the lines of a *common good*. Consider, for example, the value

³¹ There is a difficult question here about whether we have positive obligations to (say) seek out information concerning subjects about which we are currently ignorant. It seems to me that the answer is sometimes yes, sometimes no, though I will not take a stand on that question here.

³² My colleague Allan Hazlett reminds me that this sort of view, right down to the concern with apparently 'trivial' truths, can be found in W. C. K. Clifford's classic essay 'The Ethics of Belief.' For these reasons I am tempted to call the view of normativity developed here Cliffordian, at least in the sense that it suggests that epistemic normativity is ultimately grounded in moral normativity.

³³ Cf. Craig (1990), Greco (forthcoming), Weinberg (2006), Kusch (2009).

we associate with other classic examples of common goods such as *clean water*.³⁴ Given the central place of clean water in all of our lives, there is strong temptation to regard the value of clean water as intrinsic, as inherently worth possessing or acquiring, or at least as inherently worthy of our respect. From a first-person point of view, however, this does not always seem to be the case. Suppose, for example, I have my own guaranteed lifetime supply of water (perhaps a valet carries it around in jugs behind me, wherever I happen to go), more than enough to satisfy whatever thirst I might happen to have. Would any particular parcel of clean water still seem intrinsically valuable in this case—worth acquiring or possessing? It doesn't seem so, at least from my first-person point of view. And yet it looks like the nature of an intrinsic good that its goodness is a necessary feature of the thing, not the kind of property that can come and go.

That said, to appreciate the sense in which any particular parcel of water would nonetheless be worthy of our respect, suppose my imaginary valet and I are crossing a bridge over a wide expanse of clean water. What should my attitude towards this water be? By hypothesis, I am not interested in taking a drink from it; I have my jugs, after all. But while drinking this water is no goal of mine—while it is something I can find no personal value in—it seems clear enough that this stretch of water possesses a value that is worthy of my respect. Were I to dump a barrelful of sludge in the water, for example, this would clearly be something for which I would deserve blame or censure. But why's that, exactly?

Some might appeal at this point to the intrinsic value of preserving natural systems, which I would here be damaging.³⁵ But another, more straightforward answer should again strike us as plausible: namely, that other people might well need this water to satisfy their needs. Indeed, even if I have some reason to think that no one would really be harmed by the loss of this particular parcel of water, that would not seem to justify the dumping. For, given the unpredictable nature of the needs of others, and given how contamination of this sort can spread in unpredictable ways, others very well might turn out to depend on this water. And since clean water plays such an indispensable role in human well-being, we plausibly have an obligation not to pollute in this way, but rather to treat the water with due respect.

The comparison between the value of true belief and clean water is not perfect, but I think it nonetheless focuses our attention in the right way. It is not perfect, because while we could potentially flourish without clean water of any kind (perhaps we could flourish on Twin Earth, for example, with *twater*, not water), it hardly seems to be the case that we could flourish without truth of any kind.

³⁴ Goods such as *clean air* also come to mind, but focusing on water should be enough to make the point. Kusch (Ch. 3 this volume) also interestingly compares the value of true belief to the value of clean water.

³⁵ e.g. Aldo Leopold (1966: 240–2).

The comparison focuses our attention in the right way, however, because it helps us to see the way in which both belong to the category of common goods, goods that are, at least contingently, crucial to human well-being. Even though particular instances of these goods might thus not intrinsically contribute to my well-being, they should nonetheless be duly respected because of the central role that they might play in the lives of others, and perhaps even (who's to say?) in our own.

7. CONCLUSION

We can now offer a few tentative conclusions.

The first conclusion, and one that we can draw from our discussion of Goldman and Alston in particular, is that any attempt to relativize the teleological account of epistemic appraisal to 'questions of interest and importance' (or the like) faces multiple problems. For one thing, there is the problem of offering an account of what it is that makes some questions important and others unimportant, from a purely epistemic (as opposed to practical) point of view. For another thing there is the problem of how to make sense of our epistemic appraisals with respect to those questions that lie on the 'unimportant' side of the ledger (assuming such a side exists). Although in this paper I focused on the second sort of criticism, to my mind the significance of the first problem has yet to be fully appreciated, and will almost certainly prove to be the more important (and difficult) issue going forward. For example: supposing that it is true, *why* exactly is it true that counting motes of dust counts as trivial and lacking in value, from a 'purely epistemic' point of view, while finding out (say) whether one has a hand does not? Or if the hand question too counts as trivial then when, exactly, *does* one come to a question that counts as epistemically important or significant? It is hard to know how even to begin to answer these questions. And yet one often hears appeals to the epistemically 'significant' or 'important' as if these notions were well understood, or could be put to good theoretical use.

The second conclusion is that, if we follow Sosa in adopting a kind of agnosticism about the value of true belief, then we seem at a loss to explain the distinctively normative—that is, binding or action-guiding—character of our epistemic appraisals. Thus, given certain arbitrarily specified ends, we can evaluate how well someone does with respect to those ends—just in the way that Sosa describes. But it doesn't follow that a judgment that someone has failed to do well with respect to those ends is binding or action guiding in the way that a judgment that someone's belief has done poorly with respect to the truth is binding or action guiding. In order to get at the distinctively normative (as opposed to merely evaluative) character of our epistemic appraisals, we need to dig deeper.

We were then left with the following question: how *can* we make sense of the distinctively normative force of our epistemic appraisals? And according

to the proposal I sketched in the final part of the paper the best way to do this is to appreciate true belief's status as a common good (rather than, for example, as an intrinsic good, one the having of which always adds value to the life of the possessor). The way I developed this idea, moreover, was by emphasizing the fact that, as information-dependent and information-sharing creatures, we naturally—and, it seems, rightfully—depend on others as sources of information.

I will close by suggesting that if this final proposal is on the right track, moreover, then it is a mistake to suppose that the sort of normativity that characterizes our epistemic appraisals is basic or irreducible. Instead, epistemic normativity would seem to be explicable in terms of a deeper, and more obviously moral, sort of normativity: namely, the sort of normativity that derives from our obligation to help others carry out their projects and concerns (broadly understood). Although in one way this suggestion does not make the normative force of our epistemic appraisals any less mysterious—for what is the source of our non-hypothetical obligation to assist others, after all?—it does suggest that it is a mistake to try to offer an account of epistemic normativity in isolation, and without attending to what, if anything, we owe to one another from a moral point of view.³⁶

REFERENCES

- Alston, W. (1989a). 'Concepts of Epistemic Justification,' in his *Epistemic Justification*. Ithaca, NY: Cornell University Press.
- (1989b). 'The Deontological Concept of Epistemic Justification,' in his *Epistemic Justification*. Ithaca, NY: Cornell University Press.
- (2005). *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.
- Audi, R. (2001). 'Doxastic Voluntarism and the Ethics of Belief,' in M. Steup (ed.), *Knowledge, Truth, and Duty*. New York: Oxford University Press.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Clifford, W. C. K. (1877 [1999]). *The Ethics of Belief and Other Essays*. Prometheus Books, Amherst, NY: Harvard University Press.
- Craig, E. (1990). *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- Cuneo, T. (2007). *The Normative Web: An Argument for Moral Realism*. New York: Oxford University Press.
- David, M. (2001). 'Truth as an Epistemic Goal,' in M. Steup (ed.), *Knowledge, Truth, and Duty*. New York: Oxford University Press.

³⁶ Thanks to Robert Audi, Heather Battaly, Jason Baehr, Michael DePaul, John Greco, Adrian Haddock, Allan Hazlett, Martin Kusch, Alan Millar, Duncan Pritchard, Wayne Riggs, Ernest Sosa, and Fritz Warfield for helpful comments on earlier versions of this paper.

- (2005). 'Truth as the Primary Epistemic Goal,' in E. Sosa and M. Steup (eds.), *Contemporary Debates in Epistemology*. Oxford: Blackwell Publishing.
- DePaul, M. (2004). 'Truth Consequentialism, Withholding, and Proportioning Belief to the Evidence,' *Philosophical Issues*, 14: 91–112.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- (1999). *Knowledge in a Social World*. New York: Oxford University Press.
- (2002). 'The Unity of the Epistemic Virtues,' in his *Pathways to Knowledge*. New York: Oxford University Press.
- Greco, J. (forthcoming). *Achieving Knowledge*. New York: Cambridge University Press.
- Epistemic Evaluation: A Virtue-Theoretic Approach*.
- Grimm, S. (2008). 'Epistemic Goals and Epistemic Values,' *Philosophy and Phenomenological Research*, 77.
- Haack, S. (1993). *Evidence and Inquiry*. Oxford: Blackwell Publishing.
- Harman, G. (1999). 'Pragmatism and Reasons for Belief,' in his *Reasoning, Meaning, and Mind*. Oxford: Clarendon Press.
- Horwich, P. (2006). 'The Value of Truth,' *Noûs*, 40: 347–60.
- Kelly, T. (2003). 'Epistemic Rationality as Instrumental Rationality: A Critique,' *Philosophy and Phenomenological Research*, 66: 612–40.
- Kornblith, H. (2002). *Knowledge and its Place in Nature*. New York: Oxford University Press.
- Kusch, M. (2009 this volume). 'Testimony and the Value of Knowledge'
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge, MA: Cambridge University Press.
- Lehrer, K. (1990). *Theory of Knowledge*. Boulder, CO: Westview Press.
- Leopold, A. (1966). 'The Land Ethic,' in his *A Sand County Almanac*. New York: Oxford University Press.
- Lynch, M. (2004). *True to Life: Why Truth Matters*. Cambridge, MA: MIT Press.
- (2005). 'Replies to Critics,' *Philosophical Books*, 46: 331–42.
- Millar, A. (2004). *Understanding People: Normativity and Rationalizing Explanation*. New York: Oxford University Press.
- Neta, R. (2007). 'How to Naturalize Epistemology,' in D. Pritchard and V. Hendricks (eds.), *New Waves in Epistemology*. London: Palgrave Macmillan.
- Plantinga, A. (1993). *Warrant: The Current Debate*. New York: Oxford University Press.
- Quine, W. v. O. (1986). 'Reply to Morton White,' in *The Philosophy of W. V. Quine*, ed. L. Hahn and P. Schlipp. La Salle, IL: Open Court Publishing Company.
- Shah, N. (2003). 'How Truth Governs Belief,' *Philosophical Review*, 112(4): 447–82.
- and Velleman, J. D. (2005). 'Doxastic Deliberation,' *Philosophical Review*, 114(4): 497–534.
- Sosa, E. (1991). 'Intellectual Virtue in Perspective,' in his *Knowledge in Perspective: Selected Essays in Epistemology*. New York: Cambridge University Press.
- (2003). 'The Place of Truth in Epistemology,' in M. DePaul and L. Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*. New York: Oxford University Press.

- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. New York: Oxford University Press.
- Stich, S. (1990). *The Fragmentation of Reason*. Cambridge, MA: MIT Press.
- Weinberg, J. (2006). 'What's Epistemology For? The Case for Neopragmatism in Normative Metaepistemology,' in S. Hetherington (ed.), *Epistemology Futures*. New York: Oxford University Press.
- Wolterstorff, Nicholas (2005). 'Obligation, Entitlement, and Rationality,' in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*. Oxford: Blackwell Publishing.