

[Draft only]

Stephen R. Grimm
University of Montana

stephen.grimm@mso.umt.edu
www.umt.edu/phil/faculty/grimm

Epistemic Normativity

How should we make sense of our concepts of epistemic appraisal? 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 nature of this positive and negative status? What is its origin and what, exactly, is its force?¹

Among contemporary epistemologists, perhaps the most prominent way to make sense of our concepts of epistemic appraisal is in teleological terms.² On this way of looking at things, a belief earns positive marks, from an epistemic point of view, just to the extent that it helps to promote, bring about, or in some way duly respect the things with *intrinsic* epistemic value. Beliefs therefore inherit their epistemic value or status from the value of the ends that they help to bring about.

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.³ In short, the problem for the view is that our practice of epistemic appraisal 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

¹ Hilary Kornblith (2002) asks a version of our last question (and signals its elusiveness) as follows: “How is it that epistemic norms come to have their normative force? What we are looking for here is not necessarily a naturalistically acceptable answer to this question, by any acceptable answer. What we need is an account of the source of epistemic normativity that does not make a mystery of it” (p. 139).

² In addition to the figures to be discussed below, see, for example BonJour (1985, pp. 7-8), Foley (1987, ch.1), and Lehrer (1990, p. 112). For a more extensive list, see David (2001, p. 152).

³ 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.

account, I then suggest that Sosa's proposal too faces significant difficulties. I close by recommending a way of thinking about epistemic appraisal 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 Goldman, William Alston, and Michael Lynch—develop the view. Focusing on Goldman and Alston requires little by way of explanation: as two of the most influential epistemologists from the last twenty-five years, they are among the main reasons *why* the teleological account is so currently popular. Focusing on Lynch, in turn, is especially helpful because while he shares the same starting points as Goldman and Alston, he veers importantly away from their views when it comes to the crucial question of how we should think about the nature of epistemic value. As I will argue in Section 3, the Goldman-Alston version of the teleological account, on the one hand, and the Lynch version on the other, therefore point to a kind of dilemma for the teleological account. Think of epistemic value in the Goldman-Alston way, and you get one sort of intractable problem for the view; think of epistemic value in the Lynch way, and you get another sort. Not a pretty result, I will argue, for the teleological account.

Advocates of the teleological account commonly begin by recommending two claims. 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 prudential 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, Lynch not only endorse both claims, but help to show the way in which they seem to be mutually reinforcing:

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, p. 3; emphasis added)⁴

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, p. 31)⁵

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, pp. 15-16)

What I am calling the “first” claim therefore seems to be the more fundamental claim: namely, that acquiring certain epistemic goods seems to possess a kind of intrinsic value; they seem to be ends worth realizing for their own sake, and not merely for the sake of whatever further prudential ends they might help promote.⁶

⁴ This passage highlights the role of curiosity, but 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 1986 *Epistemology and Cognition*. 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, p. 98).

⁵ 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.

⁶ 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.

What I am calling the “second” claim, 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 our practical goals (for pleasure, perhaps, or for comfort) are worth realizing for their own sake, it might be less obvious that realizing our epistemic goals has the same kind of status. Indeed, it might be thought that in comparison to other goods we obviously value for their own sake, the notion of a purely *epistemic* good seems like little more than a fiction.⁷

The appeal to curiosity, I suggest, is meant to cut off just these concerns. Just as we have a range of practical desires that we naturally want to see satisfied, so too, the above passages imply, we seem to have certain distinctively *epistemic* desires, ones stemming from our natural curiosity, that we want to see satisfied.⁸

The fact that we are naturally curious about certain subjects is therefore supposed to provide support for the claim that some epistemic goods are intrinsic—worth pursuing and realizing for their own sake. But what are these 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.

⁷ According to Stephen Stich (1990, p. 131), for example, although we do value many things intrinsically—health, happiness, the welfare of our children, etc.—the truth 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.

⁸ In addition to the authors cited above, Ram Neta (forthcoming) 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” (p. 42).

Despite this initial diversity, as we'll in a moment the considered view of all three thinkers seems to be that *believing the truth* is the thing that possesses intrinsic epistemic value 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—therefore seems to have an intrinsic worth or value all its own.

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, if 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 if it does poorly with respect to the things with intrinsic epistemic value.⁹

As the following passages suggest, Goldman, Alston, and Lynch all endorse this view in very similar terms, and make plain that by their lights true belief is the thing with intrinsic epistemic value.¹⁰

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, p. 52)

⁹ This account is 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).

¹⁰ 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.” 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, p. 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, p. 128).

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, p. 28).¹¹

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. Indeed, an essential part of what makes a belief justified is precisely that it is based on grounds that make it likely to be true. If a belief’s being true wasn’t *prima facie* good, it wouldn’t make much sense to gather evidence, check and recheck calculations, or worry about whether you have good reasons for believing what you do” (2004, p. 50).

Here’s a familiar illustration of the way in which the teleological account is supposed to work. Suppose that based on your evidence it is highly likely that Jones owns a Ford. After all, Jones has been driving the same battered Ford to work for the past ten years! You therefore form a belief accordingly.¹² If it then turns out that your belief is actually false (in fact, it was really Jones’s *brother* who owned the car all along), it nonetheless seems that your belief will still earn high marks, from an epistemic point of view: it will intuitively count as justified, rational, virtuous, and so on. But why so? According to the teleological view, the reason *why* the belief earns these high marks is because it was in some sense properly responsive to the intrinsic value at issue—responsive in a way that made it (in some sense) likely that the belief would be true, even if, by chance, the truth proved elusive.

¹¹ Sharp-eyed readers will 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.

¹² To be clear: the content of the belief is “that Jones owns a Ford,” not “that it is highly likely that Jones owns a Ford.”

Although more would need to be said here about many aspects of the view,¹³ I hope that by this point 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 “promoting” or “respecting” 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

I noted earlier 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 question that begins to expose the cracks in the teleological view, is whether it is really plausible to think that *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”?

Crucially, this is where opinions begin to divide. According to Lynch,¹⁴ for example, finding out the truth with respect to *any* subject—even apparently trivial subjects like the number of motes of dust on my desk—possesses genuine, intrinsic epistemic value: it

¹³ 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.

¹⁴ Among others: Kvanvig (2003, p. 41) and Horwich (2006, p. 347) also defend this view.

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 questions concerning motes of dust and old phone numbers 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.¹⁵

Let's call this the unrestricted view of intrinsic epistemic value, according to which believing the truth with respect to any subject possesses a kind of intrinsic epistemic value.¹⁶ The reason why opinions begin to divide here is that many philosophers have found the unrestricted view of epistemic value hard to swallow, including (most notably) tried-and-true advocates of the teleological model such Goldman and Alston.

Goldman offers his own fanciful examples that apparently tell against the unrestricted view of epistemic value (1999, p. 88; see also 2002, p. 61). What is the 323rd entry in the Wichita, Kansas phone directory? Who 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 questions of this sort conspicuously *fail* to elicit our curiosity, "We can no longer suggest that higher degrees of truth possession are all that count in matters of inquiry" (2002, p. 61). Instead, on his view, we need to shift to 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, p. 61; cf. Goldman 1999, p. 89).

¹⁵ As Lynch (2004), 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 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" (p. 55).

¹⁶ Lynch (2005) 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)" (p. 331).

In his most recent book Alston too argues that an unrestricted view of intrinsic epistemic value can't be maintained.¹⁷ Since the true beliefs that we could gain from activities like memorizing phone books seem to 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, p. 32).¹⁸

Despite their initial agreement, the advocates of the teleological model we have been considering so far therefore part ways conspicuously when it comes to specifying 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.”

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, p. 61)¹⁹

¹⁷ 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” (1989, pp. 83-84).

¹⁸ 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, for example, Haack (1993, p. 199) and David (2005, p. 299). The temptation to make this move is obviously very powerful.

¹⁹ “Unitarianism” is Goldman’s term for the view that the only intrinsic epistemic value—in his words, the “cardinal” epistemic value—is true belief.

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."²⁰

What I want to argue in the following section, however, is that this difference concerning the nature of the intrinsically epistemically valuable is dramatic indeed, and that it exposes a fundamental instability 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—i.e., 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 promoted or duly respected the things with intrinsic epistemic value. But by hypothesis a true belief on such a topic would *lack* any such value.²¹ 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 18 motes. Given the sloppiness of your method, however, we can

²⁰ See, for example, Alston's 2005 discussion on pp. 30-33.

²¹ Some philosophers even insist (on a slightly different note) that a *false* belief on this topic holds no intrinsic disvalue. See, for example, Kelly (2003, pp. 624-25).

suppose that this answer really amounted to little more than a guess: you could very easily have concluded, for example, that there were a few more or a few less motes.

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 (or perhaps respect for) the truth of the matter, your belief would presumably earn low marks with respect to any type of epistemic appraisal on offer.

The problem for the restricted teleological account²² offered by Goldman and Alston, however, is to explain why this should be. Recall that according to the teleological account a belief inherits its epistemic status from its relationship to the things with intrinsic epistemic value. But if there is nothing with intrinsic epistemic value to promote (respond to, etc.) with respect to such trivial topics, then it is hard to see where this inherited epistemic value might come from. For nothing comes of nothing, as Lear once said.

The restricted teleological account therefore has a serious problem accounting for the full range of our epistemic appraisals, as opposed simply to our appraisals concerning what we might think of as “non-trivial” subjects (however that is to be understood). I will consider a possible response that Goldman and Alston might offer in a moment, but for the time being the problem for their restricted teleological account should be vivid enough that we can now see why, and despite its initial implausibility, philosophers such as Lynch have wanted to combine the teleological account of appraisal with an unrestricted view of intrinsic epistemic value.²³ For when combined with this sort of unrestricted theory of epistemic value, the teleological account has no problem making sense of our appraisals even with respect to apparently “trivial truths.” Since true beliefs concerning dust motes,

²² I.e., an account that combines the teleological view of appraisal with a restricted view of intrinsic epistemic value.

²³ I use the fudge phrase “kind of” advisedly here, for Lynch is a bit cagey about whether everything that is valuable in itself qualifies as an intrinsic value.

telephone numbers, 6th place finishers, and so on actually *are*, on the unrestricted view, intrinsically valuable from an epistemic point of view, then according to the teleological account even such trivial beliefs will be derivatively either good or bad, depending on how well they bring about the intrinsic value (i.e., truth) at issue.

This then leads us, by way of summarizing the above discussion, 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. After all, if nothing with intrinsic epistemic value is at stake, then there seems to be no sense in which a trivial belief could be derivatively either good or bad—for what would be the source of such derivative value? But it seems clear that even beliefs with respect to trivial subjects are appropriate candidates for epistemic appraisal. So combining the teleological view of appraisal with a restricted view of value seems like a nonstarter.

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 even trivial beliefs are appropriate candidates for epistemic appraisal, for such beliefs would be derivatively either good or bad to the extent that they promote (respond to etc.) 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.²⁴

²⁴ For more on the selectiveness of our sense of curiosity, see Harman (1999, p. 100) and Grimm (forthcoming).

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 nail down 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 sketched above.

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 I imagined above—as when we negatively appraise my sloppy counting of motes of dust—are not *really* problems because they are not fairly described. As I described the mote counting case, 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 have issued 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.²⁵ When a topic interests me enough to trigger a belief—in other words, when I care about it enough—it might then be thought, this bespeaks 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 is that it loses track of the fact that not just any sort of value was supposed to ground the teleological account, but rather a value that was at once both (a) intrinsic and (b) distinctively epistemic; in other words, a goal that we took to be 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 my desire to put off my work for a little while

²⁵ See Nishi Shah (2003), and Shah and David Velleman (2005).

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.

If it is not clear enough that the sort of interest in the previous examples was fundamentally non-curiosity driven, other examples are easy to find. Thus suppose you have a practical interest in boarding the next flight from Chicago to Glasgow. As such, you will clearly be interested in finding out the gate number from which the next flight departs—you will, in other words, want to believe the *truth* with respect to this subject. But it hardly seems to be the case that finding out the gate number of the next flight from Chicago to Glasgow is the sort of thing that you desire or are interested in from a strictly epistemic point, or just insofar as you are naturally curious. Instead, your interest in this subject seems entirely due to your practical interest in reaching Glasgow. It is, we might say, an instrumental interest, as opposed to an intrinsic interest.

This is a serious problem for the teleological view, however, because as we saw above it was by virtue of its relationship to things with distinctively epistemic intrinsic value that a particular belief was supposed to inherit its epistemic value or status. Substituting our practical interests—i.e., things of value and importance from a practical point of view—will therefore not get us the sort of distinctively epistemic value that (on its own terms) the account needs to make sense of our full range of epistemic appraisals.

To clarify this important point, it might help to put things in a slightly different way. Suppose we agree that insofar as we are inclined to form a belief as to whether whether p, we prefer that our belief be true rather than false. Our point here, and one which has ably been made elsewhere by Sosa (2003, pp. 157-58) and Thomas Kelly (2003, pp. 631-32), is simply that from the fact that we prefer our belief about whether p to be true rather than false it doesn't follow that a true belief about whether p is something we value for its own sake. For compare: from the fact that we prefer our tetanus shots to contain the proper

amount of the vaccine rather than a toxic amount, it surely does not follow that we value receiving proper tetanus shots, even in part, for their own sake. Instead, receiving a tetanus shot with the proper amount of vaccine is something we value only insofar as we value something else—our health—and only insofar as it looks like the shot will contribute to our health. And similarly: just because I would rather my belief about the airport gate to Glasgow be true rather than false, it does not follow that I value having a true belief on this subject just for its own sake, or insofar as I am a curious being. Instead, I seem to value having a true belief on this subject entirely because of my further practical goal—a goal, in the absence of which, a true belief on this subject would have no independent value at all.

The teleological view, at least as developed by philosophers such as Goldman and Alston, therefore has significant problems trying to make sense of our 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 offering a more subtle, and less troublesome, account of the role that value considerations play within that structure. As we’ll see in the following section, Ernest Sosa’s recent work on this topic offers just such a subtle alternative, one that should further help to clarify the nature of our concepts of epistemic appraisal.

4. Sosa’s View

A good place to begin Sosa’s account is where he himself begins, by noting that we human beings are “zestfully judgmental” across a wide range of areas: about art, literature, science, politics, sports, food, wine, and even coffee (forthcoming, p. 53).²⁶ 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

²⁶ Sosa’s account, soon to appear in book form, was first delivered as part of the John Locke Lectures at Oxford University in Trinity Term, 2005.

the values that are *fundamental* within each domain we can then appraise or assess (hence criticize) the *derivative* value of other items in the domain in terms of how well they promote, bring about, or 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 or appraise various elements of the assassin's conduct in terms of how effectively they promote 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 promote *this* fundamental goal.

Although this basic teleological framework is by now familiar to us, what focusing on these more unusual sorts of domains helps to bring out, Sosa suggests, 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 *disvalue*. 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. (forthcoming, pp. 55-56)

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 *the truth*. In this he importantly agrees with advocates of the teleological account in claiming that a particular believing is good or bad to the extent that it helps to promote, bring about, or duly respect the truth. 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 valuable absolutely, who knows? Our worry requires only that we consider truth the epistemically fundamental value, the ultimate explainer of other distinctively epistemic values" (forthcoming, p. 54).

What should we make of this view? To begin with, it has many obvious attractions. For one thing, the view easily 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 from a reliable competence on the part of the believer to realize the truth), or as “apt” or “inapt” (i.e., as realizing the truth *because of* such a competence).²⁷ 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.

5. Critique of Sosa

Despite these advantages, what I want to suggest now is that by remaining agnostic about the domain-transcendent value of true belief and by taking our epistemic evaluations to be thus “insulated,” Sosa seems to introduce a new problem—seems to, indeed, lose sight of one of the most profound aspects of our concepts of epistemic appraisal.²⁸ For notice: when we judge a belief to be unjustified or irrational, we seem to be doing more than just evaluating (in a negative way) the skill or virtuosity of the believer’s performance. In addition, we seem to be *blaming* them or *reproaching* them for believing in this way.²⁹ To judge a belief to be unjustified or irrational is therefore to judge that it *should* be given up, in some binding, and apparently culpable, sense of “should.”³⁰ As Hilary Kornblith (2002)

²⁷ Sosa has been working with this distinction for some time; see, for example, Sosa 1991. For his most recent version, see his forthcoming, ch. 2.

²⁸ At least, of what we might call our “central” concepts of appraisal: justified/unjustified, rational/irrational, and the like. He seems perfectly right about other appraisals, such as apt/inapt, etc.

²⁹ As Nicholas Wolterstorff (2005) 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” (p. 326).

³⁰ Notice that in making these claims, in suggesting that beliefs are subject to blame and censure, we do not have to suppose that belief is under our direct voluntary control. I think (along with Alston 1989b,

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” (p. 145).

Can Sosa make sense of the way in which this binding sense of “should” attaches to our central concepts of epistemic appraisal? 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.

The problem of course is that this sense of “should” is quite weak; in Kantian terms, it is the “should” of calculation rather than the stronger, binding sense of “should” that is usually associated with duty or obligation. In this weaker sense of “should,” after all, the assassin *should* use a high-powered rifle, rather than a weak pea-shooter, in order best to realize his goal of killing the target. The sense of “should” associated with our judgment that

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 judgments of blame and censure 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 judgments that a belief is (truth) unjustified carry with them the judgment that the belief should be given up, and that one is culpable for (knowingly) failing to relinquish 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 (vicious torture, for example) 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.

a particular belief is unjustified, however, seems more binding, and less a mere matter of prudence or calculation, than that.

To make better sense of this binding sense of “should” that seems to attach to our epistemic appraisals, notice that we can appraise or evaluate beliefs relative to several different goals.³¹

Suppose, by way of illustration, 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 would suddenly seem 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 goal, 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 could we be justly blamed

³¹ [[this is all very familiar to Sosa of course; he begins many of his papers with this very point]]

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 appraisals has a dramatically different character. Suppose, for example, that instead of sobering 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 likewise earn a variety of negative appraisals: thus we might say (using our artificial language of 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 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.”

Summing up, we’ve seen that we can evaluate beliefs in all different kinds of ways, relative to scores of different fundamental values and hence scores of critical domains. Given Sosa’s agnosticism about the value of truth, however, how can he make sense of the fact that the truth-perspective is in some sense the privileged perspective—the binding, non-optional perspective—when it comes to the evaluation of belief?

Here is one possible response. Since a belief that is oriented towards the truth will be much more solidly on the path to knowledge than a belief that is (say) oriented towards psychological comfort, the truth perspective will turn out to be the privileged perspective from which to evaluate beliefs simply because it is the distinctively *epistemic* perspective, in

the sense that it is the perspective relevant to *knowledge*—after all, the root of “epistemic” (*episteme*) just means knowledge.³²

The problem with this response is that the question at issue now is not what it means for a particular appraisal to count as “epistemic” but rather why appraisals made from the epistemic (i.e., truth- or knowledge-related) point of view have the particular normative, binding force they seem to have. For suppose we agree (as we should) that what it means for a particular evaluation to count as epistemic is for it to be truth- or knowledge-related. This just leads to the question: what’s so special about truth or knowledge? Why should the association with *these* properties bring with them a special, binding sense of “should”? After all, we’ve already embraced agnosticism about the value of true belief. Why shouldn’t we embrace agnosticism about the value of knowledge as well!

Sosa’s discussion of the way in which knowledge is more valuable than mere true belief³³ might suggest another way out, however. For on his view, the reason *why* knowledge is more valuable than mere true belief is because knowledge involves the manifestation of a competence (aptitude, skill) on the part of the believer, and competences are states of an agent that are able to impart value on the agent’s products. Thus, for example, just as a work of art produced by a master is more valuable or impressive than an accidental smudge that happens to reproduce the master’s strokes, so too a belief that arrives at the truth as a result of a believer’s competence or skill at realizing the truth is more valuable and impressive than a “mere” true belief—one true “just by luck.”³⁴

³² (BonJour and Sosa 2003, p. 1; though of course that’s debatable)...

³³ Commonly thought of as “the value problem.” See DePaul, Zagzebski, and Greco.

³⁴ Hence Sosa (2003) appeals to the *performance value* of a believing, “so long as the performance is high enough on the quality scale for such performances, as measured by how well the performance would provide its expected goods, if the system were properly installed” (p. 177). [[Perhaps Sosa would reply that truth, rather than comfort, is the “expected good” of the believing system, but why? If not intrinsic, then seems like a problem.]]

Is this enough to provide the truth- or knowledge-goal with the sort of extra value needed to explain why it is the privileged framework when it comes to the evaluation of belief? It doesn't look that way. If the additional, privileging value is supposed to derive from the fact that a competence is being exercised here, then the important thing to see is that this same "exercise of a competence" value can be exhibited with respect to any goal you please. Suppose, for example, I have a special competence in virtue of which my beliefs reliably realize the goal of psychological comfort. It would hardly follow that realizing this goal in a reliable way (thus achieving the comfort equivalent of knowledge) would thereby bring with it with the special sort of value associated with the binding sense of "should." Instead, everything seems to depend on whether the thing reliably (competently, skillfully, etc.) realized was worth bringing about in the first place.³⁵

6. Our Position

Let's take stock. Recall that we closed our earlier discussion of the teleological account of epistemic appraisal with a dilemma. Suppose that only truths of interest or importance are intrinsically valuable, from an epistemic point of view, and it looks like we lose our ability to explain how our epistemic appraisals apply to unimportant, "trivial" beliefs. Suppose instead that *any* truth is intrinsically valuable and it looks like we've reached absurdity; only someone really desperate would think *that*. Our discussion of Sosa's view, in turn, showed that our account of epistemic appraisal should not lose sight of the normative force of our epistemic judgments. Thus to take a belief to be justified is not

³⁵ To this Sosa might reply, finally, that not all competencies or abilities are alike, and only those that contribute to a full, flourishing life are the sort that are capable of adding extra value to the realization of a goal—are the sort, indeed, that deserve to be called *virtues*. While this seems plausible, the problem with a response of this sort—one that ties the special value of a given goal to its contribution to a flourishing life—is that it looks like the problem of trivial truths threatens to pop up yet again. Is it really the case that adding a true belief with respect to *any* subject (even the dust motes? the defunct phone numbers?) really contributes to a flourishing life? Lynch, explicitly, is happy to say so. But, as we've seen, this is an extreme position to take, and one that does not mesh well with some of Sosa's other commitments.

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 complementing when we take it to merit these positive appraisals is 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” critically leaves out.

Overall, this leaves us with at least two central elements that have proven to be extremely difficult to account for within the confines of a single theory. First, that our concepts of epistemic appraisal can be applied to the full gamut of our beliefs. Second, that our concepts of epistemic appraisal have a binding normative force, suggesting that the truth goal is not simply one goal among others but rather a goal with a special privilege or status, when it comes to the evaluation of belief.

With these two desiderata in mind, it might seem at this point that the most natural way to make sense of these two central elements is simply to bite the bullet to suppose, with Lynch and others, that any truth is (after all!) worth possessing for its own sake. *That* explains why even trivial truths are worth bringing about and promoting, and that explains why a failure to duly respect the truth brings with it criticism and reproach. But what about the absurdity!? From a first person point of view (and pace Lynch and others), it hardly seems to be the case that any truth is worth believing for its own sake, or that believing any truth adds constitutively to my well being or flourishing.³⁶

What I want to argue in this final section is that the best way to make sense our twin desiderata is in fact to accept a certain version of the thesis that any truth has a special value

³⁶ Thus Lynch (2004) argues that “Being happy is good—‘in itself’—and other things can be worth caring about deeply because they are necessary parts to it” (p. 128); Lynch then argues that believing the truth is one of the necessary or essential parts of being happy (pp. 127-28).

or worth, but from a notably different angle, and for notably different reasons, than we have considered so far. More exactly, I want to argue that the best way to make sense of these desiderata is by shifting *away* from the standard first-person question about the value of truth—wherein we ask (as Goldman, Alston, and Lynch asked) about the value of truth in terms of our own intellectual goals or well-being—and to shift instead towards a more communal or social view of the value of truth.³⁷

To appreciate how the switch from the first-person perspective to the social perspective helps to shed light on our current position, consider first the following claim from Thomas Kelly (2003):

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, pp. 624-25)

Now, some might accuse Kelly of a kind of philistinism here: that he shows himself to be unrefined in some way by failing to care about these subjects. But notice the sense in which Kelly claims not to care: from his first-person point of view there simply seems to be nothing about these topics worthy of interest or concern—nothing, that is, that would seem to contribute to his personal well-being, from a cognitive point of view or from *any* point of view.

³⁷ 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 knowledge! 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.

Rather than accuse Kelly of philistinism, I think we should concede that he is simply right about this: believing the truth with respect to such topics likely does not³⁸ contribute to any personal goals and concern he might have. What's it to him, then, if he forms a belief about whether (say) Russell was right- or left-handed by means of a random guess? And since nothing of value rides on this, from a personal, perfective point of view, it's hard to see how the belief itself could come to have derivative value or disvalue.

But now consider things from a social point of view. Suppose that Kelly does in fact opt to form a belief on one of these topics, but only by hazarding a random guess, perhaps just to pass the time during a particularly interminable wait for the doctor. Why would this belief not only deserve low marks (count as unjustified, irrational, etc.), but also earn our censure or blame, given that finding out the truth on this subject lacks any plausible value for him, cognitive or otherwise?

Looking at things from the social point of view, it seems that the deceptively simple, though for all that true, answer is that others might come along and need it; that is, that others might come along and value it. Why might *others* value it? Is there some intrinsic value that people wired differently than Kelly can appreciate that Kelly cannot? It doesn't seem that we need to go this far. Instead, all we need to recognize is that people's *practical* concerns are both enormously varied and highly unpredictable.

Consider a question such as whether Hubert Humphrey was an only child. Finding out the truth with respect to this question may not hold any value for me, may not elicit my curiosity in any way, but for Humphrey's biographer (say) getting this right will be quite important—if not in itself, or for its own sake, then only for the sake of producing an accurate account of Humphrey's life that will be accepted by publishers. But what about the

³⁸ [[Perhaps even *should not*: this would be the sense in which Sosa's "intellectual ethics" might be relevant. But I think that just shows that it is treating a quite peculiar topic.]]

really trivial topics, like the 323rd entry in the Wichita, Kansas phone directory? Or like the number of motes of dust on my desk at this very moment? Here again, to appreciate why we judge that our beliefs should be oriented to the truth even with respect to these questions, we need only bear in mind how odd and varied people's practical interests sometimes are. The Wichita phonebook fact-checker, for example, might well have a practical interest in confirming 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 are really succeeding in cutting down the number of motes on the dust on his desk, as he fondly hopes.

The basic idea we need to make sense both of the range of our epistemic appraisals but also of their normative force, then, is not that believing the truth with respect to any subject is intrinsically worthwhile or perfective, from a first-person point of view, but rather that any subject might come to have value—if only value of a practical or instrumental sort—in light of the varied and unpredictable concerns others might have. Given our place in a community of information-dependent and information-sharing creatures, we therefore plausibly have a duty to orient our beliefs towards the truth *because* we have a duty to be responsive to the practical needs of others.³⁹

Plausibly, then, when we judge a belief to earn high marks from an epistemic point of view, this is not because it is oriented towards an intrinsic good—at least if we think of an intrinsic good as one that, from the subject's point of view, is always worth realizing in the sense that it always adds constitutively to my well-being or flourishing. Instead, it seems that a better way to make sense of the value of true belief is as what we might call a *common*

³⁹ Cf. Craig, Greco, Weinberg, Kusch.

good—a good that others naturally depend on to realize their practical projects and concerns.⁴⁰

Consider, for example, the value we associate with a good such as *clean water*.⁴¹ 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 worthy of respect, as worth caring about and promoting. From a personal, perfective point of view, however, this is hard to make sense of. For one thing, is it really the case that any particular parcel of clean water possesses this value for me—that drinking it (bathing in it etc.) will contribute to my well-being? That seems like a stretch. I might, for example, 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 both to satisfy whatever thirst I might happen to have, and to provide for my own well-being. Or again, what if we were all transported overnight to Twin Earth, so it was clean twater, not clean water, that played such a central role in our lives. Would clean water still seem like an intrinsic good? It doesn't seem so, and yet it looks like the nature of an intrinsic good that its intrinsicness is a necessary feature of the good, not the kind of thing which can come and go.

And yet: suppose my imaginary valet and I are passing over a wide river, glistening with clean water. What should my attitude towards this water be? By hypothesis, I am not interested in taking a drink from this particular parcel of clean water (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—I take it that it is clear enough that this stretch of clean water is associated with a value that is worthy of my respect. Were I to dump a barrelful of sludge in

⁴⁰ Recall, for instance, U.S. Senator Barbara Boxer's criticism of Condoleezza Rice during her Secretary of State confirmation hearings in 2004. According to Boxer, Rice lacked a "proper respect for the truth" in her judgments about Iraq. Rice was, naturally, highly indignant.

⁴¹ Goods such as *clean air* also come to mind, but focusing on water should be enough to make the point. Kusch (manuscript) also interestingly compares the value of true belief to the value of clean water.

the water, for example, thus polluting it, this would clearly be something for which I would deserve blame or censure. Even though I might find no personal value in a given stretch of water clean, accordingly, that seems utterly beside the point; the water harbors a value that should be respected, regardless of whether it answers to my personal goals or concerns.

But why's that, exactly? Some⁴² might appeal to the intrinsic value of preserving natural systems, which I would here be damaging. But there is another, less controversial answer that should be as obvious as it is appealing: namely, that *other* people might well need this water to satisfy their needs. Even if I have reason to think that no one would really be harmed by the loss of this particular parcel of clear water (after all, I might point out, no one is currently downstream, and no one in the foreseeable future will be!), that would hardly seem to justify the dumping. For given the plastic and unpredictable nature of the needs of others, and given how contamination can spread in unpredictable ways, they very well *might* turn out to depend on this water. And since clean water plays such an indispensable role in our well-being, we have an obligation—to others—not to pollute in this way, but rather to treat clean water with due respect.

The comparison between the value of true belief and clean water is not perfect, but 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 (recall the Twin Earth switch), it hardly seems to be the case that we could flourish without truth of any kind. The comparison focuses our attention in the right way, however, because it does suggest that both belong to the category of common goods, goods that any of us, at any given time, might come to depend in order to realize our plastic and unpredictable goals and concerns. Even though particular instances of these goods might thus not intrinsically contribute to my well-being, and hence be worthy of pursuing for their own sake, they should nonetheless be duly

⁴² Such as, most famously, Aldo Leopold (1966, pp. 240-42).

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.

Works Cited

- Alston, William. 1989a. "Concepts of Epistemic Justification." In his *Epistemic Justification*. Ithaca: Cornell University Press.
- Alston, William. 1989b. "The Deontological Concept of Epistemic Justification." In his *Epistemic Justification*. Ithaca: Cornell University Press.
- Alston, William. 2005. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.
- Audi, Robert. 2001. "Doxastic Voluntarism and the Ethics of Belief." In *Knowledge, Truth, and Duty*. Ed. M. Steup. New York: Oxford University Press.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- David, Marian. 2001. "Truth as an Epistemic Goal." In *Knowledge, Truth, and Duty*. Ed. M. Steup. New York: Oxford University Press.
- David, Marian. 2005. "Truth as the Primary Epistemic Goal." In *Contemporary Debates in Epistemology*. Eds. M. Steup and E. Sosa. Oxford: Blackwell.
- DePaul, Michael. 2004. "Truth Consequentialism, Withholding, and Proportioning Belief to the Evidence." *Philosophical Issues* 14: 91-112.
- Foley, Richard. 1987. *The Theory of Epistemic Rationality*. Cambridge: Harvard University Press.
- Goldman, Alvin. 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.
- Goldman, Alvin. 1999. *Knowledge in a Social World*. New York: Oxford University Press.
- Goldman, Alvin. 2002. "The Unity of the Epistemic Virtues." In his *Pathways to Knowledge*. New York: Oxford University Press.
- Greco, John. Forthcoming. *Epistemic Evaluation: A Virtue-Theoretic Approach*.
- Grimm, Stephen. Forthcoming. "Epistemic Goals and Epistemic Values." *Philosophy and Phenomenological Research*. Available online at: <http://www.umt.edu/phil/faculty/grimm>
- Haack, Susan. 1993. *Evidence and Inquiry*. Oxford: Blackwell.
- Harman, Gilbert. 1999. "Pragmatism and Reasons for Belief." In his *Reasoning, Meaning, and Mind*. Oxford: Clarendon Press.
- Horwich, Paul. 2006. "The Value of Truth." *Noûs* 40: 347-60.

- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66: 612-40.
- Kornblith, Hilary. 2002. *Knowledge and its Place in Nature*. New York: Oxford University Press.
- Kusch, Martin. Manuscript. "Testimony and the Value of Knowledge."
- Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lehrer, Keith. 1990. *Theory of Knowledge*. Boulder: Westview.
- Leopold, Aldo. 1966. "The Land Ethic." In his *A Sand County Almanac*. New York: Oxford University Press.
- Lynch, Michael. 2004. *True to Life: Why Truth Matters*. Cambridge: MIT Press.
- Lynch, Michael. 2005. "Replies to Critics." *Philosophical Books* 46: 331-42.
- Neta, Ram. Forthcoming. "How to Naturalize Epistemology." In *New Waves in Epistemology*. Eds. D. Pritchard and V. Hendricks. Aldershot: Ashgate. Available online at: <http://www.unc.edu/depts/phildept/neta.html>
- Plantinga, Alvin. 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*. Eds. L. Hahn and P. Schlipp. New York: Open Court.
- Shah, Nishi. 2003. "How Truth Governs Belief." *Philosophical Review* 112: 447-82.
- Shah, Nishi, and David Velleman. 2005. "Doxastic Deliberation." *Philosophical Review* 114: 497-534.
- Sosa, Ernest. 1991. "Intellectual Virtue in Perspective." In his *Knowledge in Perspective: Selected Essays in Epistemology*. New York: Cambridge University Press.
- Sosa, Ernest. 2003. "The Place of Truth in Epistemology." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Eds. M. DePaul and L. Zagzebski. New York: Oxford University Press.
- Sosa, Ernest. Forthcoming. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. New York: Oxford University Press. Available online at: <http://homepage.mac.com/ernestsosa/menu2.html>
- Stich, Stephen. 1990. *The Fragmentation of Reason*. Cambridge: MIT Press.

Weinberg, Jonathan. 2006. "What's Epistemology For? The Case for Neopragmatism in Normative Metaepistemology." In *Epistemology Futures*. Ed. S. Hetherington. New York: Oxford University Press.

Wolterstorff, Nicholas. 2005. "Obligation, Entitlement, and Rationality." In *Contemporary Debates in Epistemology*. Eds. M. Steup and E. Sosa. Oxford: Blackwell.