

Matthew Chrisman

The Aim of Belief and the Goal of Truth: Reflections on Rosenberg¹

“What makes us cognitive beings at all is our capacity for belief, and the goal of our distinctively cognitive endeavors is truth...the basic role of justification is that of a means to truth, a more directly attainable mediating link between our subjective starting point and our objective goal” (BonJour, 1985: 7-8).

“Naturalization of epistemology does not jettison the normative...For me, normative epistemology is a branch of engineering. It is the technology of truth-seeking...it is a matter of efficacy for an ulterior end, truth” (Quine, 1986: 664).

*

“[Truth] is not what common sense would call a goal. For it is neither something we might realize we had reached, nor something to which we might get closer” (Rorty, 1995: 39).

“We know many things, and will learn more; what we will never know for certain is which of the things we believe is true. Since it is neither a visible target, nor recognizable when achieved, there is no point in calling truth a goal” (Davidson, 2000: 67).

“The upshot seems to be that no notion of truth, neither a transcendent notion of objective truth nor a minimalist notion of immanent truth, can play any determinative role at all in our epistemic activities...it begins to look as if we would do well to stop talking about truth altogether and just identify the goal of our cognitive-epistemic activities per se as justified belief” (Rosenberg 2002: 229).

1. Introduction

At the beginning of *Thinking about Knowing*, Rosenberg advertises his argumentative aim with characteristic panache: “The leading thesis developed in this book...is that knowledge is simply adequately justified belief” (2002: 1). With this, he doesn’t mean to deny that knowledge is factive: one who is committed to S’s knowing that p is committed to p’s being true. However, as he explains, “Since from any one epistemic perspective the judgements that S has done every-

¹ The original version of this essay appeared in *Self, Language, and World: Problems from Kant, Sellars, and Rosenberg*, edited by James R. O’Shea and Eric M. Rubenstein (Ridgeview Publishing Company, www.ridgeviewpublishing.com, 2010). In what follows, the text has been lightly edited and a substantive postscript added.

thing requisite to be entitled confidently to believe that *p* and that *S* has done everything requisite to establish the truth of *p* stand or fall together, a further ‘truth requirement’ is vacuous and idle. Truth may arguably be an outcome of enquiry, but it can function neither as enquiry’s goal nor as a constraining condition on any de facto epistemic policy or procedure” (ibid.: 2). Since reading this, I’ve found his idea both enticing and confusing. Ultimately, I think it is mistaken, but, as is often the case with philosophers we respect, appreciating why one of their central claims is mistaken teaches us as much or more than appreciating why many of their other claims are correct. I think a lot of what Rosenberg claims, both in that book and elsewhere is correct, so I hope it will be particularly instructive to appreciate why his idea that truth can function neither as enquiry’s goal nor as a constraining condition on any de facto epistemic policy or procedure is mistaken.

Understanding why Rosenberg claims that truth cannot function as a goal of enquiry or a constraining condition on any de facto epistemic policy or procedure requires discussing some dialectical preliminaries, after which I shall first state and then explain Rosenberg’s argument for the claim. This will involve pointing out several places where one might want to object to Rosenberg, but, in each case, I’ll argue that doing so misses his point. Finally, I shall say where and why I object to the argument. This will involve sketching a way that truth can be usefully thought of as the “aim of belief”, even if it’s not properly called a “goal of enquiry”, so that it can, pace Rosenberg, constrain epistemic policies and procedures.

2. Dialectical Preliminaries

In the course of attempting to explain epistemic normativity, it is common to treat truth as the definitive epistemic end. The idea is that there may be many different kinds of reasons to believe some proposition – e.g., prudential, moral, epistemic, etc., but what epistemology is concerned with, when it seeks to explain epistemic normativity, is what makes some consideration a good or bad epistemic reason to believe some proposition. Moreover, many have thought we might give a teleological explanation of something’s being an epistemic reason if we take truth to be the definitive epistemic end. A distinctively epistemic reason for believing some proposition is some consideration, the following of which conduces to the end of having true beliefs, rather than, say, beliefs that make you happy, are interesting, or save the world. As the quotes above from BonJour and Quine indicate, this is a prominent view; it is also popular for

the way that it makes epistemic normativity seem relatively un-mysterious because it treats epistemic reasons as a species of goal-directed reasons.²

Some philosophers deny that truth is the definitive epistemic end because they think there is no one end that can be used to distinguish epistemic normativity from other sorts of normativity. This is because they think there are several distinctively epistemic ends.³ As far as they are concerned, truth may be an epistemic end, but we have competing ends such as falsehood avoidance, epistemic responsibility, explanatory coherence, understanding, knowledge, and wisdom; and one might think these cannot be reduced to or understood in terms of the goal of truth. Depending on what we mean by “epistemic end”, I’m inclined towards something like this pluralist position. However, the crux of the issue regarding Rosenberg’s leading thesis is, I think, precisely what we should mean by “epistemic end”. So it will prove useful, even for the debate between epistemic monists and pluralists, to consider the more radical argument against the idea that truth is the definitive epistemic end. We find the beginnings of this in some suggestive passages from Rorty and Davidson (quoted above), who understand the proposal to be about our cognitive-epistemic goals. They argue, roughly, that in order for truth to be a genuine goal, we’d have to be able to tell when we achieve it, but, since we are always fallible in our beliefs and only ever have other beliefs to go on, this is impossible. Ergo truth cannot be a goal and ipso facto is not the definitive epistemic end.

3. Rosenberg’s Argument

This argument is suggestive but not rigorous. I think it is Rosenberg who has worked it out most thoroughly in ch. 6 of *Thinking about Knowing*. He claims that those who think that truth is the definitive epistemic end might be thinking of truth in one of two ways. First, they might be thinking of truth in a “realist” or enquiry-transcendent way, in that truth is not epistemically defined but something that could always in principle go beyond what we might discover. The putative goal then would be to believe that p iff p is objectively true. Second, he suggests they might be thinking of truth in a “minimalist” or enquiry-immanent way, in that truth is more or less a matter of mere disquotation. In this case, the putative goal is to believe p if and only if p .

² Other writers defending something like this idea include Foley (1987), Alston (1989), Sartwell (1992), Haack (1998: 203), Beckermann (2001), Lynch (2009, 2013), Hofmann (2005).

³ Kvanvig (2003), Riggs (2002, 2003), Pritchard (2010).

To a first approximation, then, Rosenberg's argument is this: If one is thinking of truth along enquiry-transcendent lines, then truth transcends our beliefs in that there is "no epistemically accessible truth-determinative feature of beliefs" (ibid.: 214). This is because we can never "get outside" of our belief-system to check in a non-question-begging way, whether believing for the reasons that we do brings us closer or farther away from believing what is objectively true. On the other hand, if one is thinking of truth along enquiry-immanent lines, then truth is immanent to any given belief system in that "whatever we currently believe, we hold true" (ibid.: 228). So, there is nothing more to the pursuit of the putative goal of truth about a particular proposition *p* than the pursuit of belief as to whether *p*, which means that "...the ostensible goal of immanent truth is ... a goal in name only, [since] it does not exert any constraints on the actual concrete conduct of our enquiries"(ibid.: 228–229).

Here is that argument more explicitly:

- (1) There is no epistemically accessible feature of beliefs about the world that determines whether they are objectively true. (ibid.: 214)
- (2) If there is no epistemically accessible feature of beliefs about the world that determines whether they are objectively true, then there is no way to ascertain whether our belief-forming methods result reliably in our having objectively true beliefs about the world. (ibid.)
- (3) If there is no way to ascertain whether our belief-forming methods result reliably in our having objectively true beliefs, then there is no reasonable way to evaluate the efficacy of these methods for achieving the end of our having objectively true beliefs about the world.

First lemma: There is no reasonable way to evaluate the efficacy of our belief-forming methods for achieving the end of our having objectively true beliefs about the world. (ibid.: 218)

- (4) We accept each of our beliefs to be true. (ibid.: 228)
- (5) If we accept each of our beliefs to be true, then we take each of our belief-forming methods to generate beliefs we accept to be true. (ibid.)
- (6) If we take each of our belief-forming methods to generate beliefs we accept to be true, then there is no reasonable way to evaluate the efficacy of these methods for achieving the end of our having beliefs which we accept to be true. (ibid.)

Second lemma: There is no reasonable way to evaluate the efficacy of our belief-forming methods for achieving the end of our having beliefs which we accept to be true. (ibid.)

- (7) Some end E can be a genuine goal or capable of constraining our policies or procedures only if there is a reasonable way to evaluate the efficacy of our methods for achieving E. (ibid.: 206–07)

Thus,

- (8) Neither our having objectively true beliefs about the world nor our having beliefs that we accept to be true can be a genuine goal or capable of constraining our epistemic policies or procedures. (ibid.: 229)

3.1 The Argument Explained

I shall now explain in a bit more detail each premise and why I think some reasons for objecting to several of them are misguided.

Premise (1) says that there is no epistemically accessible feature of beliefs about the world that determines whether they are objectively true. I think this should be understood as denying two things. First, it denies that there are self-evident beliefs about the world. Perhaps there are self-evident beliefs in tautologies or about the contents of one's own mind. That's debatable (Rosenberg would deny the latter sort of self-evidence for complicated Sellarsian reasons having to do with the 'myth of the given'; I'm not sure what he'd say about the former). But set these aside. When we form a belief about the world, according to premise (1), this belief never comes with an infallible marker of its truth. Second, one of the features of our beliefs may be our reasons for holding them, and, if so, (1) also denies that these reasons can ever guarantee the truth of the beliefs they support. These two denials embody a commitment to a sort of epistemic fallibilism. Rosenberg characterizes a "fallibilist understanding of the notion of justification or warrant" as "one which allows that a person may be justified or warranted in believing a proposition which is nevertheless false" (ibid: 136). Many philosophers take fallibilism to be a reflection of the in principle mind-independence of the world, or the idea that how we take things to be in the world and our reasons for doing so must always be distinguished from how the world really is. Rosenberg suggests that fallibility is "the epistemic reflection of objectivity", i.e., "that the fact that someone believes that p, or even the fact that everyone believes that p, does not imply that it is true that

p”(ibid.: 217). To fail to recognize this distinction, many think, is to commit to a crude form of idealism.

However, even while avoiding the crude idealism which rejects the in principle mind-independence of the world, other philosophers will object to fallibilism (and premise (1) of Rosenberg’s argument) based on a certain view of what our reasons for belief can be. They will say that sometimes you believe that *p* because you see that *p* or you hear that *p* or you figure out that *p* or etc., and you cannot see, hear, figure out, etc. that *p* unless *p* is true.⁴ This means that beliefs formed for these reasons are not fallible. If one’s reason for believing that *p* is that one sees that *p*, that reason does guarantee the truth of the belief.

If you’re sympathetic to this idea, then you may want to reject premise (1) of Rosenberg’s argument. However, before doing so, notice that philosophers more sympathetic to Rosenberg will worry that, even if there is a sense in which seeing, hearing, figuring out, etc. that *p* are genuine reasons for believing that *p*, they are not reasons that are epistemically accessible features of the beliefs that they support, since we often think that we have seen, heard, or figured out that *p* only later to discover that we were mistaken. In the cases where we were mistaken, it’s natural to think that our reason for believing that *p* was that we thought that we saw, heard, or figured out that *p*, but this thought turned out to be false. But why then shouldn’t that be the story about reasons in the good cases as well?

One way to diffuse this dispute in the present dialectical context is to distinguish between the reason for believing that *p* and one’s reason for believing that *p*. We can view the former as facts about the subject’s relationship to the fact that *p*, and the later as facts about what the subject takes her relationship to the ostensible fact that *p* to be. Then, if *p* is false, the reasons to believe that *p* are non-existent, since there is no fact that *p* for one to stand in some relation to. However, in such a case, one’s reasons to believe that *p* may nonetheless be things such as one’s (unfortunately false) belief that one saw, heard, or figured out that *p*.

Given this distinction, deep and perennial philosophical controversy hangs on whether we think the primary focus of normative epistemology should be the reasons for belief or one’s reasons for belief. However, without being too dismissive of the seriousness of this dispute, I think we should recognize that both proponents and opponents to the idea that truth is the definitive epistemic end, in the sense outlined above, must be committed to focusing on – in the terminology of the distinction I just made – one’s reasons for belief rather than the reasons for belief. This is because these philosophers take the claim that truth is the de-

⁴ Compare McDowell (1982, 1986), Martin (2002, 2004).

finitive epistemic end to provide not only a way of distinguishing epistemic reasons from nonepistemic reasons, but also a way of distinguishing good epistemic reasons from bad epistemic reasons. What I'm calling the reasons for belief cannot by their very nature be bad epistemic reasons, so all parties to this debate must instead be after an account of one's reasons for belief. But, if that's right, then parties to this debate cannot reject (1) on the basis of rejecting the fallibilism about reasons it embodies. And, if that's right then I don't presently see any other reason to reject (1) short of a commitment to the crude sort of idealism that rejects the in principle mind-independence of the world.

Premise (2) says that, if there is no epistemically accessible feature of beliefs about the world that determines whether they are objectively true, then there is no way to ascertain whether our belief-forming methods result reliably in our having objectively true beliefs about the world. Rosenberg's idea is not, I think, to deny that we can use one belief-forming method to check the results of another. For, surely, it is commonplace to discover that, relative to an assumed background method, some candidate belief-forming methods are more reliable than others. If, for instance, we come to doubt the veracity of the reporting on the propaganda channel and so check it against the independent news channel, we may find out that forming beliefs based on what is reported on the propaganda channel is not a very reliable way to form true beliefs, at least as judged by the reporting on the independent news channel. Instead, the idea behind (2) must be that in relying on one method to check another, we have to simply assume that the former is reliable. This means that when we check the reliability of a particular belief-forming method, we are not checking whether it is absolutely reliable but merely whether it is reliable relative to the outputs of some other method. After all, if we had instead started by doubting the veracity of the reporting on the independent news channel and then checked it against the propaganda channel, perhaps we would have reached exactly the opposite verdict. And if we checked both against our actually witnessing the events they reported on, we may discover that the independent news channel is more reliable than the propaganda channel, but only relative to the outputs of our actually witnessing the events.

What we seem to need in order to tell whether a belief-forming method is absolutely reliable is a way to compare its results directly to the objective facts. But, given (1), we are only ever able to compare the results of one belief-forming method to the results of some other belief-forming method, which means that we can never tell if we have objectively true beliefs about the facts. This is why Davidson writes, "Truths do not come with a 'mark', like the date in the corner of some photographs, which distinguishes them from falsehoods. The best we can do is test, experiment, keep an open mind. Since it is

neither a visible target, nor recognizable when achieved, there is no point in calling truth a goal” (op. cit.).

Lynch objects to Davidson’s thought here, and this is another place where you may want to object to the argument I’ve attributed to Rosenberg. Lynch suggests claims like (2) trade on the false assumption that “we can’t recognize whether a belief is true or false unless we can compare it to the naked facts”(2004: 25). Lynch thinks that the idea of comparing a belief to the naked facts is a strange sort of requirement and there is a more natural sense of recognizing when a belief is true, on which it should be a truism that we can and often do recognize when a belief is true. He writes, “The most natural interpretation of what it means to ‘recognize when a belief is true’ is that to recognize a belief as true or false is either to confirm that it is based on adequate grounds or note that it is not...In this sense of ‘recognize’ I clearly can recognize when my beliefs are true or false – it amounts to noticing whether or not they are justified” (ibid.: 26).

However, unless Lynch means to reject premise (1) by following those philosophers who claim that there are some ways of noticing a belief to be justified which guarantee to us that the belief is true, I can’t see how confirming that a belief is based on adequate grounds or noting that it is not can count as a way of “recognizing that it is true or false”. After all, even beliefs based on adequate grounds can turn out to be false.

Indeed, I think the fallibilism embodied in (1) is animated by the idea that an adequately grounded belief may nevertheless be false and a poorly grounded belief may nevertheless be true. If we accept that, then, although Lynch is of course right that when we doubt some belief we can check on its grounds or justification, this won’t provide a way to check the absolute reliability of any belief forming method. It will provide only a way to check one method against another – like checking the propaganda channel against the independent news channel. This will produce beliefs that we take to be true – in the immanent sense of truth – because we take all of our beliefs to be true; but it won’t ever put us in a position to recognize the objective truth of our beliefs. I believe this is precisely Rorty’s point in writing: “If I have concrete, specific doubts about whether one of my beliefs is true, I can resolve those doubts only by asking whether it is adequately justified – by finding and assessing additional reasons pro and con. I cannot bypass justification and confine my attention to truth: assessment of truth and assessment of justification are, when the question is about what I should believe now, the same activity” (1998: 19).

Premise (3) says that, if there is no way to ascertain whether our belief-forming methods result reliably in our having objectively true beliefs, then there is no reasonable way to evaluate the efficacy of these methods for achieving the end of our having objectively true beliefs about the world. I take this to be the least con-

trovertible premise of the argument. If there is no way to tell when we have achieved a particular end, then surely there is no way to tell which means are best for achieving it. And, likewise, if there is no way to tell whether a candidate means for achieving an end like having objectively true beliefs about the world results reliably in this end, then surely there is no way to evaluate the efficacy of that means. That's the simple idea behind premise (3). If you want to get off board here, I'm not sure what else I can say to convince you.

Premises (1)-(3) entail the first lemma. So, if I've convinced you to stay on board with them, you're committed to acknowledging that there is no reasonable way to evaluate the efficacy of candidate belief-forming methods for achieving the end of objective truth. It's worth pausing to notice how radical this conclusion will at first appear from the point of view of commonsense realism. Can we really accept that it's impossible to reasonably evaluate our attempts to hone and improve the ways we form beliefs? If that's what the first lemma of Rosenberg's argument comes to, agreeing to it may seem to involve avoiding crude idealism at the steep cost of pointless skepticism.

I think this clash with commonsense realism can be avoided. Recall the idea that Lynch and Rorty apparently agree on, viz., that trying to determine whether one of our beliefs is true or false amounts to nothing more than trying to determine whether or not it is based on adequate grounds. If we accept this, then we may think Rosenberg's talk of "objective truth" imputes a much stronger position to proponents of truth as the definitive epistemic end than they really need. Perhaps, believing truly doesn't have to be equated with infallible contact with the naked facts in order for the goal of believing truly to structure our epistemic policies and procedures. Maybe all we need is a less lofty idea of truth in interpreting Lynch's suggestion that "recogniz[ing] when my beliefs are true or false... amounts to noticing whether or not they are justified"(ibid.: 26).

It is in these terms that I think we should understand Rosenberg's claim at (4), that we accept each of our beliefs to be true. For, there is of course a sense in which we accept as true each of our beliefs that we take to be based on adequate grounds. This follows from the fact that believing that *p* entails accepting that *p* is true. So, to interpret truth in this way is, in effect, to transition to the second horn of Rosenberg's original dilemma, on which truth is disquotationally understood and so "enquiry-immanent".

In light of this, premise (5) may seem uncontroversial. It says that, if we accept each of our beliefs to be true, then we take each of our belief-forming methods to generate beliefs we accept to be true. This links a platitude about beliefs to a correlative conceptual fact about our belief-forming methods. If it's part of believing that we take each of our beliefs to be true, then it's part of something's being one of our belief-forming methods that we take it to generate something,

i. e. beliefs, which we take to be true. In spite of this, I think (5) is problematic; however, I'll postpone discussion of my worry until later.

Premise (6) says that, if we take each of our belief-forming methods to generate beliefs we accept to be true, then there is no reasonable way to evaluate the efficacy of these methods for achieving the end of truth. I think the idea is that, since it's part of their being belief-forming methods that these methods lead to our accepting some proposition as true, there is no meaningful sense in which we can evaluate them as more or less reliable for achieving this end. Insofar as they are methods that form beliefs, they will always achieve this end, so there is no way that they can serve as constraints on actual epistemic policies and procedures. As Rosenberg puts the point, "...the ostensible goal of immanent truth is a goal in name only...[For] nothing can count as a reason to believe that we have failed to reach or realize the putative goal"(ibid.: 228).

However, premises (4)-(6) entail the second lemma. So, if we accept them, we're committed to acknowledging that there is no reasonable way to evaluate the efficacy of our belief-forming methods for achieving the end of our having beliefs which we accept to be true.

That leaves premise (7), which says that some end E can be a genuine goal or capable of constraining our policies or procedures only if there is a reasonable way to evaluate the efficacy of our methods for achieving E. As I said above, I think that a lot rides on how we're to think of the ostensible end of truth, which means that a lot rides on premise (7). I suspect it is false as stated. However, I'll postpone criticism until I've finished explaining the argument. For, even if we're thinking of the end of truth as a goal, premise (7) then states a necessary condition on being a genuine goal which some have objected to.

In his review of Rosenberg 2002, Fantl objects to the idea that for something to be a genuine goal, it must be possible to check whether we've achieved it. He writes "Why buy into the premise that the only thing that can count as [an epistemic] goal is something that can constrain our epistemic practices or that we can confirm we've reached?"(2007: 231). As a counterexample, Fantl asks us to consider his goal that his children thrive after his death. He writes, "Regardless of whether my life goes worse if that goal is not achieved, there is a clear sense in which I have it as a goal and a clear sense in which I have failed to achieve that goal if they do not thrive after I am dead"(ibid.). However, I think this misses the point of premise (7) in the argument.

The point is not to deny that some of our goals are for ends that we will not, as a matter of fact, be able to confirm whether they have been achieved. The point is rather that the allegedly definitively epistemic goal of truth is one for which it would be either impossible to tell that we have achieved it or impossible to tell that we have not achieved it. It's quite clear how Fantl could tell whether

he has achieved his goal that his children thrive after he is dead. He'd just have to wait a sufficient amount of time after his death and then check on how his children are doing. To be sure, that isn't humanly possible, but it's pretty clear what it would take for us to check whether the goal had been achieved. By contrast, opponents of the putative goal of truth think it's unclear what it would take for us to check that we had achieved the goal of having objectively true beliefs and utterly pointless to check that we had achieved the goal of having beliefs that we take to be true.

Nevertheless, we might still worry about premise (7) for related reasons. For it may seem that it is enough, in order for something to be one of my goals, that I take it that some means are better than others. Consider, for example, the ostensible goal of getting into heaven, which, for the sake of argument, suppose we cannot even in principle ascertain whether we have achieved it since heaven is pure bliss and devoid of unified consciousnesses. In this case, (7) implies that getting into heaven cannot be a genuine goal; however, it seems clearly possible to have this as one of our goals and let it prescribe particular conducts, e.g., praying, going to church, etc. insofar as we take these practices to be encouraged by having that goal.

I think Rosenberg's initial answer to this objection comes in the distinction he draws between a motive and a goal:

... a motive is whatever in fact moves someone, and there's no in-principle limit to the sorts of mental goings-on that might in fact give rise to activities that we'd recognize as enquiry – a desire for money or fame or respect...In this sense, a desire for objective truth might indeed motivate enquiry. But it doesn't follow that objective truth can function as the goal of enquiry, and that becomes clear when we observe that, unlike such other potential motivating desires as those for money or fame or respect, a desire for objective truth is one that we can't ever determine has been satisfied.(ibid.: 219fn.)

If a motive is whatever in fact moves one to act, then what is a goal? The causal-psychological description of a motive suggests that the contrasting notion of a goal could be understood in terms of the distinction between motivating and validating practical reasons. A motivating reason is part of the best psychological explanation of why an agent acts as she does, and a validating reason is whatever would validate a particular action (whether or not the agent actually performs it or performs it for that reason). On this reading, Rosenberg's claim is that a desire for truth can provide a motivating reason for particular actions, but truth cannot provide a validating reason, and thus is not a goal, because it's impossible to tell whether it has been achieved. So, a goal is distinguished from a mere motive in that it validates, i.e. provides a validating reason for, the relevant action.

But in order to evaluate this idea and to see how it supports (7), we need an account of “validating reasons” in the context of the question of whether truth is the distinctive epistemic goal. What are the relevant actions that may or may not be validated in this case? To begin, Rosenberg proposes the following:

The specification or formulation of a goal...is characteristically the first step in a process of means-ends reasoning. In traditional terms, it yields the major premise of a practical syllogism:

- (P1) I/We shall achieve E
 (P2) The only/best means for achieving E are M
 (P3) So, I/we shall adopt M. (ibid.: 206)

The idea here is that since the practical syllogism is the traditional formulation of instrumental reasoning per se, we can define the sense of “goal” that is relevant to the argument by means of the functional role of goals in practical syllogisms. With respect to the activity of enquiry, Rosenberg is specific about what he thinks it would take for truth to be the relevant goal. He writes:

What we need, to make sense of the idea that true belief is the goal of enquiry...is an instantiation of this general practical syllogistic form which establishes a connection between a commitment to that goal and the actual conduct of enquiry; that is, which shows how our having the goal of truth can structure our actual concrete cognitive-epistemic practices (ibid.: 207).

The general idea seems to be that some end E can provide a validating reason for doing M only if we can produce an instantiation of the practical syllogistic form which establishes a connection between doing M and achieving E. Whether this supports (7), however, will depend on what it means to produce an instantiation of the practical syllogistic form which establishes a connection between doing M and achieving E. For in one sense, it has to be trivial that, if some of our belief forming methods are indeed objectively reliable, then it is logically, physically, and epistemically possible that we can produce an instantiation of the practical syllogistic form where the minor premise is objectively true. And this instantiation would, in a sense, establish a connection between these methods and attaining the goal of objective truth. The problem with this, however, is that we cannot ever tell that we have produced this. That is to say, even if we have done this, since the fallibilism of (1) committed us to the idea that the truth of our beliefs transcends our reasons for believing them, we cannot distinguish a validating practical syllogism from indefinitely many others which do not validate any particular action. So, in stronger sense, we would not have established the necessary connection unless there is a reasonably way to evaluate whether

candidate means for achieving x are more or less likely to succeed. So, I think Rosenberg means that, in order for an end E to be a genuine goal or capable of constraining policies or procedures, it must be possible for us to produce an instantiation of the practical syllogistic form which we can tell to establish a connection between doing particular conducts M and achieving E . If that's right, then (7) appears to successfully state a necessary condition on being a goal. Moreover, with the first and second lemmas, it entails Rosenberg's leading negative thesis, viz. (8): Neither our having objectively true beliefs about the world nor our having beliefs that we accept to be true can be a genuine goal or capable of constraining our epistemic policies or procedures.

This completes my explanation of the argument suggested by Davidson and Rorty and developed most thoroughly by Rosenberg against the idea that truth is an epistemic end. I've attempted to motivate each of its premises, though I've already said that I think there's a problem with premises (5) and (7). I want to turn now to spelling out these problems. I'll start with premise (5) because the objection I have to it is more easily stated, but the real insight Rosenberg's argument promises, is in forcing us to rethink his conception of epistemic ends as goals, which is basically what funds my objection to premise (7).

4. Two Objections

Recall that premise (5) is a conditional linking the platitude about accepting each of our beliefs to be true to the idea that we take each of our belief-forming methods to generate beliefs we accept to be true. Here it is again:

- (5) If we accept each of our beliefs to be true, then we take each of our belief-forming methods to generate beliefs we accept to be true.

The problem I see with this stems from an ambiguity in the term "our belief-forming methods". Premise (5) is trivially true if this means "each of our actually and currently deployed belief-forming methods"; for part of what it is to be an actually and currently deployed belief-forming method is to be a method which generates a mental state in the person who deploys it that this person currently accepts as true. However, if we separate belief-forming methods across members of a population or across time-slices of an individual person, we can generate a different reading of the term "our belief-forming methods" that makes premise (5) much less plausible. For instance, if that term means "the candidate belief-forming methods that might be deployed by one of us" then surely we don't take each of our belief-forming methods to generate beliefs

that we all accept as true. After all, some of us form beliefs by watching the propaganda channel, and, as we have already seen, this generates beliefs that those of us who are viewers of the independent news channel will not accept as true, and vice versa. Likewise, a single individual may have formed beliefs in the past by watching the propaganda channel, which she then accepted as true, but, after she has become a devotee of the independent news channel, she will not take it that forming beliefs by watching the propaganda channel is a reliable way to generate beliefs that she now accepts to be true.

So far, all this suggests is that Rosenberg needs the reading of “our belief-forming methods” on which it means “each of our actually and currently deployed belief-forming methods”, in order for premise (5) of his argument to be plausible. However, recall that (5) is a premise in Rosenberg’s argument along the second horn of his original dilemma for the conclusion that immanent truth, i.e. believing p iff p , cannot be a genuine goal of our cognitive-epistemic practices. Someone who thinks that immanent truth is a goal of our cognitive-epistemic practices might reasonably insist that they don’t mean for us to use this goal to evaluate the efficacy of our actually and currently deployed belief-forming methods but rather for us to use this goal to evaluate candidate belief-forming methods, such as those used by others, those we used to use in the past, and those we might use in the future. And it is not in general true that the belief-forming methods used by others, our past selves, or our future possible selves are ones that generate beliefs that we now accept as true.

What this means is that although Rosenberg is of course right that we accept each of our own beliefs as true, he’s wrong that this implies that there is no reasonable way for us use the ostensible goal of immanent truth to evaluate candidate belief-forming methods. To be sure, as long as we have accepted premises (1) and (2), such evaluations will always transpire relative to an assumed background of belief-forming methods and so will never get at the absolute reliability of candidate belief-forming methods for achieving the end of objective truth. However, the second horn of Rosenberg’s dilemma is not concerned with such objective truth. And it does seem that, relative to a background method, a method like the propaganda channel can be judged to be less reliable than another method, such as the independent news channel, for forming true beliefs. If that’s right, however, then Rosenberg’s premise (5) is acceptable only on a reading which his opponents could reasonably reject.

I turn now to premise (7). In effect, I object to premise (7) because of the univocal conception of ends on which it rests. As we have seen, it rests on a conception of ends as goals that can serve as the first premise in a practical syllogism. I think this is misguided in the present dialectical context because it assumes that one who thinks that some end, such as truth, may help to distinguish epistemic

reasons from other sorts of reasons must be conceiving of epistemic reasons as fundamentally reasons for some action, i.e. practical reasons for some epistemic-cognitive conduct.

I doubt, however, that the primary objects of epistemic evaluations are epistemic-cognitive actions or conducts. Rather the primary objects of epistemic evaluations seem to me to be states of belief. After all, we evaluate whether someone's belief that *p* is justified in determining whether she knows that *p*, and normative epistemology is concerned with what count as good and bad reasons for a particular kind of mental state: belief. To see those evaluations as deriving from evaluations of some implicit cognitive conducts strikes me as backwards.

At the very least, this suggests a worrying disconnect between premise (7) of Rosenberg's argument and attempts to use the putative goal of truth to distinguish good epistemic reasons for belief from other sorts of reasons. Yet, it may not seem like a very deep worry if we could come up with a cognitive action or conduct which is conceptually very close to the cognitive state of belief and the reasons we may have for such states. In a way, this is what Rosenberg suggests as part of his own positive theory. He appeals to Sellars' (1963) "forward-looking" and "proceduralist" conception of justification, according to which

justification is itself in the first instance a cognitive-epistemic activity – giving reasons, citing evidence, enumerating grounds – and being in a position to engage in such activities can and often will itself be a consequence of having engaged in other, preparatory and enabling, cognitive-epistemic conducts – seeking reasons, evidence, or grounds for what one in fact believes, however one in fact has come to believe it. (ibid.: 212)

This is ultimately why he thinks, "the relevant executable epistemic conducts can be specified only generically, as whatever is necessary to put one into a position to justify one's beliefs if they are legitimately challenged" (ibid., p. 213).

However, there's a much simpler way to reconcile the fact that the primary object of epistemic evaluations are cognitive states of belief with a teleological conception of such evaluation. This is by appealing to something active that is conceptually very close to the state of belief: the formation of beliefs. Rosenberg would probably object that although forming beliefs is undoubtedly active rather than stative, and surely also a central element of our cognitive activities, it is not properly speaking what he calls an "executable conduct" – it is active but not something we freely choose to do (at least not in normal cases), rather like digesting one's dinner or falling asleep. What that means is that, while forming beliefs – like digesting dinner or falling asleep – is, broadly speaking, something that we do, it cannot cogently serve in the conclusion of a practical syllogism meant to represent (even if only schematically and ideally) the way truth is the goal of belief-formation. And that means that it cannot comprise the episte-

mic policies and procedures that are supposed to be constrained by the ostensible goal of truth.

That objection, however, exemplifies the way in which I think premise (7) is misguided. Why must we think that reasons for forming beliefs must ultimately be capable of being characterized as practical reasons that justify performing some “executable conduct”? Rosenberg might respond: “Well, how else do you propose to characterize them?” So, let me quickly sketch a provisional proposal.

Indeed, I think it is Sellars who provides the distinction crucial for developing an alternative. In a different context, he distinguishes between “rules of criticism” and “rules of action”, i.e. between rules that articulate how something ought to be and rules that articulate what agents ought to do. (1969: 508) For instance, a car ought to be disposed to start when the key is turned and my mechanic ought to do what is necessary to make my car be disposed to start when the key is turned. These ought-statements express two different kinds of norms. Although, as this example indicates, they can be logically related: that my car ought to be disposed to start when the key is turned implies *ceteris paribus* and where possible that my mechanic ought to do whatever is necessary to make my car disposed to start when the key is turned.

What I like about Sellars’ distinction in the present context is that it opens up conceptual space for a mixed account of normative epistemic principles, according to which the genuine executable conducts of agents are subject to rules of action, while cognitive states and all of the other mere activities and performances that make up our cognitive lives are only subject to rules of criticism.⁵ And this means that we can ask what beliefs we ought to form, and understand this question in terms of what beliefs we have reasons to form, without assuming that these reasons must be reasons for performing some action conceived of on the model of the practical syllogism validating some “executable conduct”. Instead, they are conceived of as a distinctive kind of reasons: reasons for being in a particular state.

It is because of the availability of this sort of account of reasons for belief and the fact that proponents of the idea that truth is the definitive epistemic end or goal have been concerned to develop an account of epistemic normativity that distinguishes reasons for belief from other sorts of reasons that I think premise (7) of Rosenberg’s argument is misguided. The way that it presupposes that

⁵ I developed a similar idea in more detail in my (2008), whose central thesis is that ‘ought-to-believe’s should be thought of as rules of criticism; this allowed me to diffuse the debate between doxastic involuntarists, who think belief formation is not voluntary, and epistemic deontologists, who think that beliefs can be the proper subjects of epistemic obligations.

all reasons are essentially practical reasons obscures the fact that epistemic reasons for belief are not practical reasons but something fundamentally different: reasons for being a particular way.

Someone might want to defend Rosenberg here by arguing that whether or not epistemic reasons for forming beliefs are fundamentally different from practical reasons, those philosophers who seek to distinguish epistemic reasons from other sorts of reasons in terms of the putatively definitive epistemic end of truth are committed to following Rosenberg in his conception of the relevant reasons as being essentially practical reasons. This is because an end just is something we pursue, and we pursue it by performing particular actions, i.e. engaging in particular “executable conducts”.

To begin to see why this is wrong and why premise (7) is false, notice that the logical form of (7) appears to be a conditional that embeds a disjunction in its antecedent. The two disjuncts are

- (a) Some end E can be a genuine goal,
- and
- (b) Some end E can be capable of constraining our policies or procedures.

The premise then says, in effect, that if (a) or (b) is true, then there will be a reasonable way to evaluate the efficacy of our methods of achieving E. However, to put my objection bluntly, I think Rosenberg has illicitly assumed that (a) and (b) come to the same thing. That is, what a goal is, on his conception, is an end capable of constraining our policies and procedures; and an end capable of constraining our policies and procedures just is, on his conception, a goal. However, I want to argue that, even if a goal must be capable of constraining our policies and procedures, it’s not the case that all ends, which are capable of constraining our policies and procedures are goals. If that’s right, then it’s possible that there are ways for ends to constrain our policies and procedures other than by structuring a practical syllogism that validates particular actions or concrete conducts.

In particular, I want to suggest that there are two ways an end can be related to some active performance. First, as we have seen, an end can be a regulative goal of that performance; regulative goals can be used to evaluate various means in terms of their efficacy at attaining the goal. Second, an end may also be a constitutive aim of a particular performance; appeal to constitutive aims can be used to determine whether or not one is engaged in some specific type of performance. For example, the legal liability of a doctor when her patient dies may depend on whether the doctor was treating illness (however poorly) or murdering someone (however skilfully). Relative to the goals, respectively, of

preserving life and of terminating life, we can evaluate how skilful the doctor's conduct was; this is an evaluation of a performance in terms of its regulative goal. But skill is beside the point in deciding whether the doctor should be charged with malpractice or murder. In order to answer this question, presumably we need to form an opinion about which type of performance she was engaged in, which we can do if we have an opinion about what its constitutive aim was. This is an evaluation of a performance in terms of its constitutive aim.

Importantly, as I just stated the distinction between regulative goals and constitutive aims, it applies first and foremost to ways of evaluating performances. Now, some performances may be genuine actions, i.e. what Rosenberg thinks of as the executable conducts that can be referred to in the conclusion of a practical syllogism, but, as we've already seen, not all active performances are executable conducts, some are mere performances, like digesting one's dinner or falling asleep. And, as I've already said, forming beliefs seems to be a mere performance rather than a genuine action (at least in most cases). That may mean that it is incoherent to speak of the regulative goal of this performance. In forming beliefs, one is not, strictly speaking, trying to attain some goal. This is because one isn't trying to do anything; forming beliefs isn't (typically) an "executable conduct". However, that doesn't mean that this performance doesn't have a constitutive aim. Digesting one's dinner and falling asleep are both end-directed performances, even though they are not goal-promoting actions. We can evaluate whether some performance is one of these specific types of performances in terms of whether it is aimed at the relevant end. Likewise, I think this opens a new route to understanding the claim that truth is an epistemic end, usable in an account of what distinguishes good epistemic reasons for belief from other sorts of reasons. We could treat truth as a constitutive aim of the performance of forming beliefs rather than as the regulative goal of some action or cognitive conduct. The idea would be to say that what constitutes a performance's being the performance of forming beliefs is (at least in part) that it is aimed at the truth. But what does that mean?

I admit that it's not entirely clear to me, but in what follows, I'll provisionally suggest an answer that I think is plausible and undermines premise (7) of Rosenberg's argument. In order to do so, it's helpful to note two things about rules of criticism. First, they seem to divide into intrinsic and extrinsic rules. Second, their logical grammar seems to allow for an 'in order to E' operator. Let me explain.

The rules of criticism applying to something seem to divide into those which are intrinsic to what that thing is and those which are extrinsic to what that thing is. For example, the rules governing the bishop in chess seem to include: (i) The bishop ought to be moved only diagonally, and (ii) The bishop ought not to be

sacrificed for a pawn. The first of these ‘ought-to-be’s is intrinsic in that it partially defines what it is for something to count as a bishop; if this rule is broken with respect to a particular thing, it calls into question whether that thing is a bishop. By contrast, the second of these ‘ought-to-be’s is extrinsic in that it is not partially definitive of what it is for something to be a bishop; if this rule is broken with respect to a particular thing, it does not call into question whether that thing is a bishop, only whether the owner of the bishop is a good chess player.

Moreover, all rules – both ‘ought-to-do’s and ‘ought-to-be’s – seem to admit of modification with an ‘in order to E’ operator. That is to say that, if they aren’t already so modified, we can make the full logical form of statements deploying the concept ‘ought’ more explicit by adding ‘in order to E’, where E is a variable that can be filled in by various ends. We can refer to this as the end-relativity of ‘ought’.⁶ For example, it may be the case that one ought to give to charity in order to develop the character trait of being charitable or in order to assuage one’s guilty conscience, or perhaps both. Depending on what the relevant end is, we get logically distinct readings of “A ought to give to charity”. Likewise with rules of criticism: It may be the case that one ought to be charitable in order to help the world or in order to impress one’s mother, or perhaps both. Again, we get logically distinct readings of “A ought to be charitable” depending on the relevant end. The end-relativity of ‘ought’ shows up with the rules of criticism governing belief-formation as well. For example, it may be the case that someone ought to be disposed to form the belief that she can jump across certain crevices, in order to believe what’s true, but it may also be the case that she ought to be disposed to form this belief in order to have the confidence to jump far enough to make it, or perhaps both.

In light of the distinction between intrinsic and extrinsic rules of criticism and the general end-relativity of ‘ought’, I think we can make better sense of the idea that truth is the constitutive aim of the formation of beliefs. To a first approximation, the idea would be to say some ‘ought-to-be’s governing beliefs are intrinsic to belief and belief formation while others are extrinsic, and we can tell the difference by the end to which the oughts are explicitly or implicitly relative. More specifically, it’s the ‘ought-to-believe’s which are implicitly or explicitly modified by ‘in order to believe truly’ that are intrinsic rules of criticism for beliefs, whereas those ‘ought-to-believe’s that are explicitly or implicitly modified by some other end are extrinsic rules of criticism for beliefs.

⁶ Compare Finlay (2009).

So, for example, consider the claim “You ought to believe that you are reading this paper.” One thing this could mean is that you ought to believe that you are reading this paper in order to believe truly. Another thing this could mean is that you ought to believe that you are reading this paper in order to make me happy. The current interpretation of the idea that truth is the definitive epistemic end would treat the former but not the latter as constitutive rules of criticism applying to that belief. That is to say that violations call into question something’s being a belief.

That’s suggestive, but it won’t work as it stands. A false belief that *p* is one that would violate the rule “*S* ought to believe that *p*, in order to believe truly”; so, if this rule really is constitutive of something’s counting as a belief that *p*, it’d be impossible to believe falsely without calling into question the status of one’s belief, as a belief. A helpful modification of the proposal can be gleaned from Velleman, who argues that an important difference between belief and desire is that for a propositional attitude to be a belief it must (at least) involve regarding its propositional object as true, whereas for a propositional attitude to be a desire it must (at least) involve regarding its propositional object as to be made true. But, although this so-called “difference in direction of fit” distinguishes belief from desire, Velleman argues that more is needed to distinguish belief from other cognitive attitudes such as imagining or assuming. This is because these too involve regarding a proposition as true, just not seriously or in earnest. However, according to him “what distinguishes a proposition’s being believed from its being assumed or imagined is the spirit in which it is regarded as true, whether tentatively or hypothetically, as in the case of assumption; fancifully, as in the case of imagination; or seriously, as in the case of belief” (2000: 183). This suggests the following modification of the proposal from the previous paragraph. We could say that it’s the ‘ought-to-believe’s which are implicitly or explicitly modified by ‘in order to have a propositional attitude we seriously regard as true’ that are intrinsic rules of criticism for beliefs, whereas those ‘ought-to-believe’s that are explicitly or implicitly modified by some other end are extrinsic rules of criticism for beliefs. This, I think, provides a more plausible interpretation of the idea that truth is the definitive epistemic end.

As I said above, I myself suspect that truth is not the definitive epistemic end because I’m sympathetic to the pluralist position that there are other distinctively epistemic ends, which are not reducible to the end of truth. However, the present point is that by understanding the idea as a claim about constitutive aims rather than regulative goals we seem to block any conflation of (a) and (b), the two disjuncts of the antecedent of premise (7) of Rosenberg’s argument. And by doing so, we open up conceptual space for the possibility that truth is an end capable of constraining epistemic policies and procedures even if it is

not a goal, in the sense of being a possible first premise in a practical syllogism which turns on the efficacy of certain actions vis-à-vis that goal. It does so just like the end of moving diagonally constrains the movements of the bishop in chess. If the rule of criticism, “Bishops ought to be moved diagonally in order to conform to the rules of chess,” is violated, that violation calls into question whether the relevant thing counts as a bishop. Likewise, if the rule of criticism, “Beliefs ought to be held in order to have a propositional attitude we seriously regard as true” is violated, that violation calls into question whether the relevant thing counts as a belief. And this provides a sense in which truth is an end constraining epistemic policies and procedures but not a genuine goal, which is capable of serving cogently as the first premise of a practical syllogism that validates particular cognitive actions or executable epistemic conducts.

5. Conclusion

So, in the end, although Rosenberg may be right that objective truth cannot be a goal constraining our actual epistemic policies and practices, there are, as far as I can tell, two cogent responses defenders of the idea that truth is a distinctively epistemic end may make to his argument. First, they may reject his imputation to them of a conception of truth as enquiry-transcendent and yet reject premise (5) along the enquiry-immanent horn of his original dilemma. I challenged this premise based on different readings of “our belief-forming methods”. Second, they may reject the conception of epistemic ends as goals capable of serving as the first premise of a valid practical syllogism, which was the idea behind (7). We saw this to be questionable based on possible conflation of the first and second disjuncts of the antecedent of (7); and I’ve tentatively sketched an alternative way of conceiving of truth as an epistemic end that doesn’t involve seeing it as a regulative goal but as a constitutive aim of belief.

6. Postscript

Since I wrote this paper, an interesting literature on “normativism” about belief has flourished.⁷ As I understand it, this is the view that an essential part of what it is for something to be a belief is for it to be governed by certain kinds of norms;

⁷ See McHugh and Whiting (2014) for a useful survey of the recent literature and some of the key issues.

or similarly, being governed by an identifiable kind of norm is partially constitutive of being a belief. This is usually conceived as a central part of the explanation of which sorts of normative evaluations of a belief count as “epistemic”.

The most widely discussed form of normativism seems to be the view that belief is essentially governed by a “truth norm”. It’s controversial how exactly to state this norm, but the basic idea is that an essential part of what it is for a state to be a belief is for its truth value to determine whether the state is correct/incorrect, right/wrong, permissible/impermissible.⁸ This idea is, I think, related to but not exactly the same as the idea I sketched in this paper that truth might be usefully thought of as a constitutive aim of the performance of forming beliefs. So, I suspect it will be useful to identify similarities and differences.

We can come at this by considering some criticisms of both ideas. The critical argument from Rorty-Davidson-Rosenberg that I attempt to reconstruct in the main text of the paper is similar to one of the prominent criticisms of belief normativism due to Glüer and Wikforss (2009, 2013a). Working from the assumptions that norms are supposed to be prescriptive or at least be capable of guiding us, they argue that a truth-norm could never do that directly, since we’d have to already believe *p* in order to follow a rule such as “You ought to believe *p*, only if *p*!” Although supported differently, this is very much in line with Rosenberg’s conclusion that a supposed goal of truth cannot “play a determinative role at all in our epistemic activities”(2002: 229) because there is no way to evaluate candidate means of achieving it, which don’t depend on one’s having already done everything one could know to do to achieved it.

What should the truth normativist say in response to Glüer and Wikforss? I take it they should (and sometimes do) say something similar to what I said in response to Rosenberg: There is more than one way that the end of truth could be related to belief. Being what Rosenberg refers to as a regulative goal that could fund a prescription or rule capable of being followed in concrete executable conduct is not the only way for truth to be an end. It might also be a constitutive aim, an end the relation to which constitutes something as the formation of a belief rather than something else.

In the main text of the paper, I tried to make good on this suggestion by appealing to a distinction Sellars makes in connection to linguistic rules between “rules of action” and “rules of criticism”. The former say how someone ought to act, while the latter say how something ought to be (which implies, *ceteris paribus* and in some complicated way, that someone ought to act in some way). For

⁸ Defended in one form or another by Millar (1994), Shah and Velleman (2005), Boghossian (2008), Littlejohn (2012), Wedgwood (2002), and Whiting (2010).

example, that an eighteen-month-old child ought to be disposed, *ceteris paribus*, to say “red” when queried about the color of red things is not plausibly thought to be a rule that guides that child’s linguistic behavior. And this rule of criticism is only indirectly prescriptive for the behavior of the child’s caregivers in that it implies that, when they are trying to teach the child words for colors, they ought *ceteris paribus* to take whatever concrete means are at their disposal to inculcate, encourage, and stiffen the child’s disposition to say “red” when queried about the color of red things.

As far as that goes, however, a rule of criticism needn’t be constitutive of anything. The previous example illustrates this: the fact that one ought to be disposed to say “red” when queried about the color of red things isn’t constitutive of being an eighteen-month-old child.

Nevertheless, I think we could get a version of the normativist position out of rules of criticism by accepting that the following rule of criticism is constitutive of belief: beliefs ought to be true. If you find a propositional attitude that is not true, then either it is liable to a kind of criticism as not true or it is not a belief. A normativist might say that this is a necessary condition on something’s being a belief. Adhering to Sellars’s distinction then provides materials for a response to Glüer and Wikforss: We needn’t think this rule (the “truth norm”) is a rule that directly guides people’s actions, in particular the agent’s actions that result her having a belief. For it could be a constitutive rule of criticism, only related to rules of action in some complicated way. Hence, I think rules of criticism, when they are thought to be constitutive of belief, provide a way to work out the normativist position about belief in a way that is immune to the objection that these norms cannot guide action.

This is an interesting position. I’m sometimes attracted to it, but I think it is stronger than the view I sketched in the main text of the paper and I harbor worries about it. As I indicated in the main text of the paper, I’m sympathetic with a pluralist view, which would find room for some sort of plurality of definitive epistemic ends. So even if it is essential to something’s being a belief that it ought to be true, I doubt that is the only thing that determines which evaluations of belief are epistemic. If that’s right, then truth normativism would only be part of the story about which norms constitutively govern belief.

But more to the point, I now suspect there are several ways philosophers and ordinary people use the term ‘belief’, which makes me skeptical that there is one unified kind of thing properly thought of as belief, whose essence could be normative in the way suggested in the previous paragraph. Just to name a few: there are the rationally formed conclusions of careful deliberation, the pieces of common sense inculcated by nature and culture, the things we accept on faith, the mere opinions (or *doxa*) about some matter of the non-experts, the

information tacitly possessed by normal people, the ways our pets conceptualize and track reality, the representational states attributed to intentional systems (including corporations, animals, computers, and ecosystems) as part of a particular kind of explanation and prediction.

Perhaps there is some common essence to all of the things on this list, but I doubt it's being governed by anything usefully thought of as a norm, even a "rule of criticism". The normativist could, of course, maintain that not all of these are properly called "belief" or even just stipulatively limit the focus of his account of belief to a subset of this list.

I suspect, however, that there is a way the things on this list are unified in their relation to the end of truth, even if this isn't normative in the sense that seems to be assumed by the normativist. This comes into sharper relief when we remember that it's a metaphor when we say that belief aims at truth. Since belief is stative rather than active, on the standard categorization deriving from Aristotle, Vendler, and Ryle,⁹ it doesn't even fall into the metaphysical category of things whose telicity (or relation to an end) helps to define them. So, if anything belief aims at truth in a sense derivative of the fact that belief formation is a telic process. The proximal end of this process is having a belief, but plausibly one isn't engaged in belief-formation unless the result when this process completes is something that can be evaluated as correct/incorrect in terms of being true/false. Does that mean that all such states ought to be true (in the way of a "rule of criticism")? If we accept panoply of kinds of belief alluded to above, the affirmative answer to this question seems implausible. There may be many systems for which it is usefully to describe them functionally as believing *p* but where it would be rather odd to think that anyone ought, even *ceteris paribus*, to do anything to ensure that such beliefs are true.

If we shift attention from the beliefs of our colleagues, friends, and children to those of corporations, animals, computers, and ecosystems, even weak normativism developed with the concept of a rule of criticism begins to seem to strong. If we shift attention in this way, the force of any normative epistemic evaluations also weakens considerably. Hence, the lesson I draw is that some kind of constitutive norm might apply to the beliefs of people whose cognitive efforts we evaluate as epistemically justified/unjustified, good/bad, permissible/impermissible, etc. But if so that is not constitutive of the relevant states being beliefs but rather of the sort of epistemic community within which the normative evaluation of beliefs and believers makes sense.

⁹ See Chrisman (2012) for more discussion of this categorization and why it's important for normative epistemology that belief is a state rather than a process.

References

- Alston, W. (1989). *Epistemic Justification*. Ithaca (NY): Cornell University Press.
- Beckermann, A. (2001). Wissen and Wahre Meinung. In: W. Lenzen (ed.), *Das weite Spektrum der Analytischen Philosophie: Festschrift für Franz von Kutschera*. Berlin: Walter de Gruyter.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge (MA): Harvard University Press.
- Chrisman, M. (2008). Ought to Believe. In: *Journal of Philosophy* 105: 346–370.
- Chrisman, M. (2012). The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions. In: *Journal of Philosophy* 109: 588–612
- Davidson, D. (2000). Truth Rehabilitated. In: R. Brandom (ed.), *Rorty and His Critics*. Malden (MA): Blackwell Publishing.
- Fantl, J. (2007). Review of Thinking About Knowing. In: *Philosophy and Phenomenological Research* 17.
- Finlay, S. (2009). Oughts and Ends. In: *Philosophical Studies* 143: 315–340.
- Foley, R. (1987). *The Theory of Rationality*. Cambridge (MA): Harvard University Press.
- Haack, S. (1998). *Manifesto of a Passionate Moderate*. Chicago: University of Chicago Press.
- Hofmann, F. (2005). Epistemic Means and Ends: In Defense of Some Sartwellian Insights. In: *Synthese* 146: 357–369.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lynch, M. (2009). Truth, Value and Epistemic Expressivism. In: *Philosophy and Phenomenological Research* 79: 76–97.
- Lynch, M. (2013). The Truth of Values and the Values of Truth. In: D. Pritchard (ed.), *Epistemic Value*. Oxford; New York: Oxford University Press.
- McDowell, J. (1982). Criteria, Defeasibility, and Knowledge. In: *Proceedings of the British Academy* 68: 455–479.
- McDowell, J. (1986). Singular Thought and the Extent of Inner Space. In J. McDowell and P. Petit (eds.), *Subject, Thought, and Context*. Oxford: Clarendon Press.
- Martin, M. G. F. (2002). The Transparency of Experience. In: *Mind and Language* 17: 476–425.
- Martin, M. G. F. (2004). The Limits of Self-Awareness. In: *Philosophical Studies* 120: 37–89.
- Pritchard, D. (2010). The Value of Knowledge. In: A. Haddock, A. Millar and D. Pritchard, *The Nature and Value of Knowledge*. Oxford, New York: Oxford University Press.
- Quine, W. V. O. (1986). Reply to Morton White. In: P. Schilpp and L. Hahn (eds.), *The Philosophy of W. V. Quine*. La Salle, Ill: Open Court.
- Riggs, W. (2003). Balancing Our Epistemic Ends. In: *Noûs* 37: 342–352.
- Riggs, W. (2002). Beyond Truth and Falsehood: The Real Value of Knowing That P. In: *Philosophical Studies* 107: 87–108.
- Rorty, R. (1995). Is Truth a Goal of Enquiry? Davidson Vs. Wright. In: *Philosophical Quarterly* 45: 281–300.
- Rosenberg, J. (2002). *Thinking About Knowing*. New York: Oxford University Press.
- Sartwell, C. (1992). Why Knowledge Is Merely True Belief. In: *Journal of Philosophy* 89: 167–180.
- Sellars, W. (1969). Language as Thought and as Communication. In: *Philosophy and Phenomenological Research* 29: 506–527.
- Velleman, D. J. (2000). *The Possibility of Practical Reason*. Oxford: Oxford University Press.

