

[Note: This is a draft in progress; please cite only with permission.]

The Epistemic Point of View¹

J. ADAM CARTER

✓ (K ((
j.a.carter@ed.ac.uk

5.1 Introduction

The topic that will be pursued here is that of epistemic value generally, and more specifically, epistemic evaluations. The aim will be to investigate what makes something more or less good from an (point of view—that is, from the point of view from which it is our epistemic purposes that matter.

Epistemologists tend to align themselves with one of two particular ways of thinking about the epistemic point of view. I'll refer to the first way as (epistemic value monism), and the second as . What divides veritists and pluralists in epistemology is the role truth should be thought to play in epistemic evaluations. According to the veritist, truth is the fundamental epistemic aim. This is not the claim that nothing but truth is epistemically valuable; rather, the idea is that something is more or less good epistemically in virtue of its being the sort of thing that promotes truth (or avoids error). Accordingly, then, the veritist says that:

$U((((())))$): The point of view from which what matters is maximizing truth and avoiding error.

Veritism has a venerable tradition in contemporary epistemology; as a thesis, it is both simple and forceful. The pluralist view, on the other hand, is comparatively less elegant. According to the pluralist, truth is but one of several, perhaps many, epistemic aims. And accordingly, the

¹ Special thanks to Ernest Sosa, Pavel Davydov, Josh Orozco, Matt Benton and Alex Jackson for comments on an earlier draft of this chapter at the epistemology dissertation seminar at Rutgers University. Thanks also to Trent Dougherty, Alvin

pluralist will tell us that the epistemic point of view is one from which truth is not the only thing that matters, and that truth does not exhaust the possible sources of epistemic value.

Given veritism's comparative explanatory power along with its simplicity, it should be no surprise that this view has been a favourite axiology in mainstream epistemology. Veritism has the advantage of parsimony for one thing, and this parsimony makes possible a straightforward way of explaining why, for example, one belief but not another counts as justified or why one trait is an intellectual virtue and another a vice.

But is veritism true? On the basis of (might we think that veritism is true and pluralism is false? This will be the central issue I'll be exploring here. In what follows, I will first make precise what it is we should take the veritist to be arguing for and then consider how the most viable veritist line holds up against a variety of challenges in recent literature. Although no strategy for defending veritism will ultimately turn out promising, I argue that a core idea of veritism can nonetheless be preserved within a particular kind of pluralism—one that upholds the truth goal alongside a distinct goal that governs agency.

5.2 The strong case for veritism

It is common to say that beliefs aim at truth. Perhaps equally common, though, is the thought that the beliefs we hold should be rational or justified. Do we want our beliefs to be rational and justified simply because we want to have rational and justified beliefs? More likely, it seems, we want rational and justified beliefs because they are markers of truth, which is what we really want. If that is right, then a neat and simple picture seems to emerge: we think our beliefs are better or worse to hold on the basis of their propensity to guide us to the truth, and we also think of beliefs themselves that they at truth. It's quite natural then to connect these observations and suppose that the latter fact must play a role in explaining the former: that it's

((((that we inquire, as believers,(i) -41 Q 2(g) -8 (u (t-7 (g) 2(g) -81 0 Tm 1 Tf [(a) -8 ()-8 (7m) .

some cognitive faculties, belief forming processes, methods of inquiry, etc. better than others. It's easy, then, to see how the thought that truth is our fundamental epistemic aim is compelling: beliefs aim at truth, and epistemic agents, qua believers, inherit this aim, which in part explains why we hold as epistemic values such things as rationality, justification, memory, and those traits we call intellectual virtues; they are epistemically valuable because truth is our goal and they help get us to it.

5.3 The normative and teleological accounts of truth-directedness

If the fact that belief aims at truth is supposed to play a role in explaining why truth is the fundamental epistemic value, as the strong case for veritism I just outlined suggests, then we should insist on two arguments: one which clarifies and defends the claim that beliefs aim at truth, and another that establishes how it is that the truth-directedness of belief would stand to support the veritist's claim that truth is the goal of inquiry.

Let's look first at the claim that beliefs aim at truth. As Engel (2004) points out, this claim is at best metaphorical: 'Obviously beliefs do not 'aim' at anything by themselves, they do not contain little archers trying to hit the target of truth with their arrows.' (Engel 2004: 77). So what does it mean then to say that beliefs aim at truth?

According to Engel, the claim $\langle \langle \langle \rangle \rangle \rangle$ is best understood as one expressing the idea that truth is the norm of belief. Beliefs, he says, have $\langle \langle \langle \rangle \rangle \rangle$ properties, which means that 'if certain properties are not instantiated, they are incorrect or wrong.' (Engel 2004: 81). An important normative property of beliefs, Engel claims, is that they are correct only if the proposition believed is true. He calls this the $\langle \langle \langle \rangle \rangle \rangle$:

$U \langle \langle \langle \rangle \rangle \rangle$ (Engel): For all P, one ought to believe that P only if P. (2004: 82)

Along with Engel, Lynch (2008) also endorses the idea that truth is the norm of belief. According to Lynch:

$U \langle \langle \langle \rangle \rangle \rangle$ (Lynch): It is prima facie correct to believe $\langle p \rangle$ if and only if $\langle p \rangle$ is true. (Lynch 2008: 4)

It's important here to note that both Engel and Lynch take the truth norm to be more than just an observaton about beliefs, but moreover a constitutive fact about the very concept of belief. Says Engel:

...a subject cannot ascribe to herself a belief unless she recognizes that belief is constrained by the norm of truth...in this respect the connection between the very concept of belief and the epistemic norm of truth is conceptual or constitutive." (Engel 2004: 83)

This idea is shared by Velleman (2000) who proposes that "the concept of a belief is just the concept of an attitude for which there is such a thing as correctness or incorrectness, consisting in truth or falsity. (Velleman 2000: 16) Lynch, drawing from Velleman, follows suit and says that the truth norm "states a constitutive fact about belief" (Lynch 2008: 7).

If Engel, Velleman and Lynch are right that part of what it is for something to even be a belief is that it counts as correct (or one you ought to hold) only if true, then the claim that belief aims at truth is a claim; that beliefs are truth-directed is part of what makes them count as beliefs. Now we may ask: does this imply anything about whether truth is the fundamental epistemic aim? It certainly seems plausible that it would; after all, the truth norm is taken to be a normative constraint on belief.

Lynch argues for just such an implication; he thinks that by engaging in the practice of inquiry you thereby assent to recognizing the value that governs the practice, and for the practice of inquiry—which consists in forming beliefs—that value is truth. And so for Lynch, then, you are committed to recognizing truth as the value that governs inquiry just by engaging in inquiry at all.

Lynch's move here commits what Engel takes to be the mistake of supposing that:

in order to be responsive to the norm of truth, believers must recognize this norm and reflect upon it... (Engel 2004: 88)

Engel adds:

...believing that p entails believing that p is true, but it does not involve any attitude towards the proposition 'P is true.' It only involves an attitude towards P." (Engel 2004: 88).

Engel here distinguishes between the teleological account of truth-directedness and the normative account of truth directedness. Whereas the normative account says only that it is constitutive of the concept of belief that it is governed by the truth norm, the teleological account identifies believing with both '(a) the conscious recognition of the basic norm of truth and (b) the intention to respect and maintain this norm in the formation of one's beliefs.' (Engel 2004: 84) We can see, then, that veritism will be an implication of the claim that belief aims at truth, then, only if the teleological account of truth-directedness is correct as opposed to the weaker normative account of truth-directedness. But, as it seems, the teleological account is correct only if responsiveness to the truth norm requires reflecting upon it, or having some attitude toward it. While these requirements would appear to hold for intentional acts such as

'accepting' or 'judging,' it remains that 'the believer of an individual belief does not consider the aims of inquiry when he comes to believe a given proposition.' (Engel 2004: 86) To note, this is entirely compatible with the thought that the truth norm is a constitutive part of what belief is.

Says Engel:

Being able, at least tacitly, to recognize the epistemic norms of beliefs is certainly a necessary condition for having the concept of a belief. But in order to be a believer, one does not need to attend reflectively to these norms. (2004: 88)

Consequently, we see how it is that the strong case for veritism as a thesis motivated by the truth-directedness of belief loses its purchase power entirely. Engel puts this point aptly:

So the fact that belief aims at truth (in the normative sense) is one thing, and the fact that inquiry has truth as its goal is another thing... the believer in an individual belief is not someone who contemplates the general goal of aiming at truth and acts upon it. (2004: 88)

So if veritism is true and truth is the fundamental aim of inquiry, this wouldn't simply be because belief aims at truth. If veritism is true, it must be defended some other way.

5.4 A different case for veritism

I've argued to this point that if veritism is true, it won't be because beliefs aim at truth. Now I want to defend a more simple point, which is that if veritism is true, it couldn't be a (truth. How might one defend veritism as a conceptual truth? A tempting strategy would be to take as a starting point something like Sosa's (2007) suggestion that 'Evaluation is distinctively epistemic when it is concerned with truth' (Sosa 2007: 78) and then try to show that veritism is straightforwardly implied by this apparent truism. Consider the claim:

E (: Something is evaluated epistemically iff what matters for the purposes of evaluation is just getting to the truth.

The conceptual line states a rather intuitive idea, and veritism is directly implied by it; for truth will be the fundamental epistemic aim if indeed getting to the truth is the goal against which epistemic evaluations are made. But can the veritist really defend veritism simply by offering up a claim about what is constitutive of epistemic evaluations? It should be clear that such a strategy would be ineffective and certainly incapable of gaining any traction against the opponent who endorses the pluralist view. The reason this sort of conceptual line defense of veritism would be ineffective is that 'epistemic' functions here in a premise that would serve to support veritism rather than pluralism only if 'epistemic' is ((((it would mean only if veritism (the conclusion) rather than pluralism is true. More plainly: the veritist

could never persuade the pluralist with an argument whose premise is one that could be accepted only by one who already accepts the veritist's conclusion. The veritist, therefore, cannot argue in a way that assumes that the epistemic point of view is just the point of view from which getting to the truth is what matters. After all, how we should characterize the epistemic point of view is specifically a point of contention between the two positions, and so the pluralist could not in principle be moved to conviction by this sort of reasoning.

The veritist must therefore not assume in her argument that 'epistemic' means something the pluralist could not accept. What we may insist on from the veritist then is some non-questionbegging argument that establishes veritism to correctly characterize () ((((from which the considerations we weigh when making evaluations are just those that matter in light of our properly epistemic (understood non-committally) purposes. This is what we should insist on from the pluralist as well.

It should now be a bit clearer what a successful defense of veritism would have to look like. And this said, I think William Alston (2005) provides a helpful way of understanding how the epistemic point of view might be characterized, non-questionbeggingly, along veritist lines. Here's Alston:

We evaluate something when we dub it good, bad, or indifferent for some purpose or from some point of view.... 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. (Alston 2005: 28)

Importantly, what Alston says here allows that the veritist and pluralist can agree that the epistemic point of view is the point of view from which what matters is determined by () while at the same time disagreeing about what these purposes actually are.² And so here we see wherein the substantive debate lies. Accordingly, then, are in a position to assess whether veritism is correct by considering what our epistemic purposes in fact are. If they just are the purposes we have insofar as we're after the truth, then the veritist has got things right, and the pluralist goes home emptyhanded.

² Here Alston explains why the purposes that count as 'epistemic' are those directed at acquiring true rather than false beliefs. Because he thinks that a viewpoint from which things are evaluated as good or bad will be an 'epistemic' viewpoint in virtue of its being a viewpoint from which things are good or bad () (((((, it will follow that the epistemic point of view is viewpoint from which things are more or less good insofar as they promote the purpose of attaining true beliefs rather than false beliefs—the veritist's claim.

5.5 Kvanvig, David and the “Knowledge Objection” to Veritism

Our aim then is to figure out whether the veritist is right about what our epistemic purposes are. If some of our epistemic purposes are ones we have not only insofar as we’re after truth, then veritism is in trouble. Against this background, I want to introduce what might seem at first to be a simple and obvious objection to veritism along these lines, but which turns out to be quite complicated and one that will deserve our sustained attention. Let’s call this the *Knowledge Objection* to veritism; it is one that has surfaced most notably in an exchange between David (2001) and Kvanvig (2005). What I call the *Knowledge Objection* takes as its initial premise the conclusion we reached at the end of the previous section:

- (1) If veritism is true, then our epistemic purposes are just the purposes we have insofar as we’re after the truth (i.e. maximizing true beliefs, avoiding false beliefs). (Sec 5.4)

The next step in the objection is to offer up something that seems uncontroversial:

- (2) Our epistemic purposes consist in part in purposes we have insofar as we’re after knowledge.

But what is jointly implied by (1) and (2) is:

- (3) Therefore, veritism is false.

This simple argument carries with it some compelling force.³ After all, why should we think—as the veritist does—that epistemic purposes must be restricted to those purposes we have insofar as we’re after something true? It should be obvious here that the Meno Assumption lends some credence to Premise (2). As Kvanvig

³ Suppose, for example, that Bill truly believes that his wife is having an affair. Though he’s already got a true belief, he nonetheless desires to know whether she is having an affair. Of Bill’s situation, we might be inclined to think that the purposes Bill has insofar as he’s after knowledge are purposes over and above, and distinct from, the purposes he has insofar as he’s after what he’s already got (true belief). And if the latter purposes matter from an epistemic point of view, then why shouldn’t the former? To put the question differently:

(2009) puts it:

When asked, “why, from a purely cognitive point of view, do you value knowledge?” we are inclined to answer...that we want to be correct, but not merely by accident, as what happens when one has a merely true belief. (Kvanvig 2009 : 7)

Insofar as we want knowledge rather than mere true belief, then, we seem to have epistemic purposes distinct from the purposes we’d have if only wanting, as Kvanvig says, to ‘be correct...by accident.’ (2009 : 7)

In sum, then: Premise (1) in the () is an articulation of the substantive () point the veritist must be making in order to avoid begging the question against the pluralist, and (2) captures the seemingly uncontroversial idea that our epistemic purposes, at least some times, are those we have insofar as we want to (), and not just merely truly believe. But (1) and (2) jointly entail the denial of veritism. Could objecting to a thesis as widely held as veritism really be this easy?

5.5.1 Marian David’s “Circularity” Response to the Knowledge Objection

For obvious reasons, the knowledge objection is one to which the veritist had better have a quick and ready reply. According to Marian David (2005)⁴ the best way to meet this objection is to call a circularity foul on proponents of the objection. Here’s David:

Although knowledge is certainly no less desirable than true belief, the knowledge-goal is at a disadvantage here because it does not fit into this picture in any helpful manner. Invoking the knowledge-goal would insert the concept of knowledge right into the specification of the goal, which would then no longer provide an independent anchor for understanding epistemic concepts. In particular, any attempt to understand justification relative to the knowledge-goal would invert the explanatory direction and would make the whole approach circular and entirely unilluminating (David 2005: 153-154).

David’s beef with the knowledge-goal objection to veritism, then, is that by placing knowledge as the epistemic goal, we would be forced to give a theory of justification that explains justification by appealing to knowledge. For example, the reliabilist—if trading the truth goal for the knowledge-goal—would have to say that a belief is justified insofar as it is produced by a process that reliably gets to knowledge. But if this is how the reliabilist understands justification, then her theory of knowledge as true belief that reliably gets to knowledge will be circular. David sums up the point:

...[That] Knowledge was supposed to be explained in terms of justification and not the other way round...does not mean that it is wrong in general to talk of knowledge as a goal, nor does it mean that epistemologists do not desire to have knowledge. However, it does mean that it is bad epistemology to

⁴ Philip Ebert has noted in discussion a similar concern with the knowledge goal.

invoke the knowledge-goal as part of the theory of knowledge because it is quite useless for theoretical purposes: The knowledge-goal has no theoretical role to play in the theory of knowledge (David 2005: 153-154)

David's point about circularity is initially compelling: a theory of knowledge that appeals to knowledge in that theory's definition of knowledge would be circular. But David's reply here becomes far less compelling under closer scrutiny.

5.5.2 Kvanvig's Response to David

Now Kvanvig (2005) has offered up a rather provoking response to David's reasoning. First, Kvanvig points out that David has assumed that epistemology is the theory of knowledge (a point he makes earlier), and secondly, David takes it that 'epistemic evaluation... is most naturally understood along broadly teleological lines, as evaluating beliefs relative to the standard or goal...' (Kvanvig 2005: 10). Such a picture leaves, as Kvanvig notes, justification as the "only epistemic concept in the theory of knowledge" (Kvanvig 2005:10)—that is, within epistemology. With these assumptions in place, the only candidates for epistemic goals other than truth will be justification or knowledge itself. Between knowledge and justification, knowledge stands to be the obvious competitor. But because knowledge as the goal was thought to lead to vicious circularity within a definition of knowledge, only the truth goal remains as a candidate. And so by process-of-elimination, veritism must be true.

Taking this to be the David's reasoning, Kvanvig argues, at the outset, that:

The form of argument here bears scrutiny since the particular argument in question relies on the false assumption that epistemology is the theory of knowledge. (Kvanvig 2005: 10)

Kvanvig is I think correct to take note that it was this assumption that made it seem inevitable that either knowledge or truth must be the epistemic goal, the position from which David argued against knowledge on grounds of circularity and in favor of truth. Because Kvanvig thinks that the epistemology-as-theory-of-knowledge assumption is mistaken, he says that we will need a generalization of David's argument in order to conclude that truth is the primary epistemic goal. (2005: 10) Here's Kvanvig:

Such a generalization could begin by allowing that any purely theoretical cognitive success is of value and hence a suitable epistemic goal. But for each such goal, it has no theoretical role to play within the project of theorizing about it, for such an account would make the theory 'circular and entirely unilluminating.' Moreover it would have to be claimed, any goal that would render a theory circular and unilluminating in this way cannot be an epistemic goal, leaving truth as the only standing candidate for that role. So, on this [David's] view, only values that fall outside the

domain of cognitive successes can legitimately be cited when explaining epistemic success.⁵ (Kvanvig 2005: 10)

In a nutshell, then, Kvanvig is saying that once the epistemic goal—that is, whatever it is that matters primarily from the epistemic point of view—is taken to be the goal of epistemology insofar as epistemology is a theory of knowledge, we would be in a position to (as David does) discount the knowledge-goal as circular and be (by process of elimination) veritists only at the expense of having to grant both that (i) nothing of epistemic value can be an epistemic goal, and (ii) that only values that fall outside the domain of cognitive successes can explain epistemic successes.

If Kvanvig is right, then David is paying simply too heavy a price to pay in order to prop truth up as the epistemic goal rather than knowledge. First off, it's no good being stuck with the counterintuitive consequence that nothing of epistemic value could be an epistemic goal.⁶

Similarly problematic is the other implication of David's view that Kvanvig draws attention to, which is that only values that fall outside the domain of cognitive successes can explain epistemic successes. This implication would be problematic for two reasons: first, central proponents of veritism (e.g. Goldman) commonly articulate the epistemic goal twice over: as the goal of truth, but also as the goal of maximizing true beliefs. After all, when the veritist tells us which traits count as virtues, she points to whichever ones best lead us to form true beliefs. But unlike 'truth' (full stop), 'true beliefs' a variety of cognitive successes, and straightforwardly so. So David's circularity response to the knowledge objection to veritism would have the bizarre consequence that veritists cannot explain justification in terms of true belief—a consequence that would mean no reliabilist could be a veritist.

5.5.3 What does the knowledge objection tell us about veritism?

We must resist the temptation to infer, from the previous discussion, that the knowledge

⁵ (?(

⁶ For one thing, the thought that nothing of epistemic value can be an epistemic goal would be at odds with David's conception of epistemic evaluation "along broadly teleological lines." For if we think of epistemic evaluation along teleological lines, we would have to understand, for example, the epistemic value justification has as epistemic value—value justification has as a means to promoting some epistemic end. But it is a prominent view in value theory (i.e. Korsgaard (1983); (Rabinowicz (2000) that the instrumental epistemic value justification should be thought to have could only be understood in the context of its promoting whatever is epistemically valuable, though epistemically valuable nonetheless. That is, whatever it is that is epistemically valuable for its own sake, and not for the sake of something else. That said, it's a problematic outcome, at best, for David to be forced to deny that truth—his epistemic goal—is epistemically valuable. It's not clear in fact how to even make sense of such a claim.

objection is sufficient for undermining veritism; all we've shown is that the knowledge objection can't be rebutted via the circularity argument David gives. Even more specifically, though, what was shown was that the veritist can't defend against the objection that knowledge is an epistemic goal over and above truth by arguing in a way that assumes that epistemology is nothing more than the theory of knowledge. It is appropriate to ask, then, whether the veritist could reply to the knowledge objection by appealing to some presumed circularity at issue with the knowledge-goal while restricting, as David did, epistemology to simply the ((.

5.4 Ecumenical veritism

As Hillary Kornblith (2002) has observed, veritists indeed do have the tendency to, much as David has seemed to, 'reduce epistemology to the theory of propositional knowing.' (Axtell & Carter 2008: 3). But that's just an accidental feature of many veritist views; they don't have to think of epistemology this way. Let's see how veritism might fare alongside a wider conception of what epistemology is. Toward this end, let's think very basically: if we take metaphysics to be concerned predominantly with how things , then epistemology, in an equally broad sense, might plausibly be thought of as concerned with our representations of how things are. Wayne Riggs (2004) offers a helpful way to think about our plight in the broad sense of representers of the world:

I believe that we humans, as cognitive beings, ultimately strive for a representation of the world and our place in it that is comprehensive, accurate and intelligible to us. And the cognitive apparatus we have to achieve this is fallible, limited, and largely dependent on the contributions of others. (Riggs 2004: 1)

Accordingly, it should seem that whatever else epistemology is a theory of, it should be a theory that concerns both what it is we strive for in our representations of the world and what it is that makes us more or less apt to success. Now Sosa (2007) might be inclined to insist on caution here by drawing the distinction as he has between the theory of knowledge and intellectual ethics and supposing that we are here conflating the two. He writes:

Accordingly, we do well to distinguish between two parts of epistemology: (a) theory of knowledge, and (b) intellectual ethics. The latter concerns evaluation and norms pertinent to intellectual matters generally, with sensitivity to the full span of intellectual values. It is therefore a much broader discipline than a theory of knowledge focused on the nature, conditions, and extent of human knowledge. (Sosa 2007: 89)

I think Sosa is right to point out, as he has, that some evaluations about beliefs are based on considerations independent from whether the beliefs are accurate, and I'll agree that such evaluations fall outside of the area of epistemic normativity; these then will not be epistemic evaluations. But we can grant this while, at the same time, granting that epistemic normativity ranges over inquiry in a broader sense than David, for example, has assumed.

Hookway (2000, 2006) for one takes epistemology—and epistemic normativity—to range over matters concerning how we are able to carry out our inquiries in a well-regulated manner. This core idea has been adopted not only by epistemologists with pragmatist leanings like Hookway, it should be noted, but also by philosophers steeped in the areteic tradition, such as Hursthouse (1999) and Zagzebski (1996).⁸

There's no reason the veritist can't take on board this wider purview. Let's call veritism free from the 'epistemology-as-a-theory-of-propositional-knowledge' assumption (). The question becomes: are our (ecumenical) purposes—purposes we have insofar as we aim to () to best () — just the purposes we have insofar as we're after the truth (i.e. maximizing true beliefs, avoiding false beliefs)?

Enter again now the knowledge objection. Ecumenical veritism would indeed avoid the problems associated with the initial "theory of knowledge" assumption with which Kvanvig took issue, and it would do so at the cost of no longer being able to frame just truth (knowledge as the only candidate epistemic goals. (It was this dichotomy, recall, that led David to think truth was the only candidate goal left standing). To be clear, though, the cost of giving up this dichotomized picture of the candidate epistemic goals is that the veritist is no longer in a position to conclude that the circularity of the knowledge goal would veritism (by process of elimination). It remains open, though, for the ecumenical veritist to try to defend itself against the knowledge objection by arguing that the knowledge goal is circular.

Perhaps the best case that could be made to this effect would be to argue that () what epistemology is a theory of, it is at least in part concerned with what knowledge is. And so if, as the objection goes, knowledge is an epistemic goal distinct from truth, then David's worry seems to reemerge: justification would be defined in terms of knowledge, and so knowledge can be defined in terms of justification only at pain of circularity.

This strategy is less ambitious than the one David himself gave us in that it doesn't attempt to establish that veritism is —it aims only to establish that the knowledge objection fails to show that veritism is false. It's also more plausible than David's in that it doesn't rely on

⁸ This agent-centered conception of epistemology, with a focus on epistemic practice as well as epistemic ends, chimes closely with the view endorsed by Zagzebski (1996). While Zagzebski holds the rather controversial view that epistemic success is a subspecies of moral success, she maintains the rather comparatively compelling view that the measure of epistemic success—that is, the end against which things are measured as good or bad epistemically—is the end consisting in epistemic flourishing, or alternatively, the 'epistemic good life.' Hursthouse (1999) holds a similar view here. On this note, one might think, for example, of Aristotle or Socrates as archetypal examples of epistemic flourishing. On such a picture, Hookway's end of "being a good inquirer" or Zagzebski's end of epistemic flourishing would leave us saying that for some given item of evaluation O, O has epistemic value to the extent that it promotes the fundamentally valuable epistemic end of being the sort of inquirer that we take Aristotle or Socrates to have been.

an overly restricted purview of epistemology. Despite these relative improvements David's circularity charge has with this new-and-improved ecumenical view of epistemology's subject matter we've equipped it with, the circularity charge will still miss the mark. This is because it (its mark only if the pluralist proponent of the knowledge objection could explain those particular concepts in terms of which knowledge is defined (i.e. justification) only by appealing to knowledge—that is, only if the pluralist, when defining those concepts that knowledge is defined in terms of (let's just use justification for simplicity), is committed to what I'll call the ((

:

✓ ((: If theory T recognizes E1, E2...En as epistemic goals, then whatever account of justification T gives must be one that explains justification in terms of E1, E2...En.

If a pluralist offers the knowledge objection to veritism while at the same time committed to the ((, then her theory of knowledge would indeed be circular; knowledge (an epistemic goal would function indispensably within an account of justification, and thus will be imbedded inevitably and perniciously within a theory of knowledge.

But there is no reason that someone who offers the knowledge objection must also be committed to the ((. More specifically, nothing about the mere countenancing of multiple epistemic goals would count for or against what theory we should hold about which goals are to be appealed to when evaluating something positively or negatively epistemically.⁹ It is thus open to the pluralist to explain what makes a belief justified in terms of some epistemic aim distinct from knowledge while at the same time allowing that knowledge is an epistemic aim in its own right.

5.5 The 'collapsing' defense

I think it is only fair that we consider here briefly a line of response the veritist has available against the knowledge objection that relies on the assumption that epistemology is the theory of knowledge (on the supposition that the knowledge-goal is circular. A remaining idea that the veritist could appeal to would be as follows: the purposes we have insofar as we are

⁹ I want to draw here from a point Riggs (2004) makes when discussing the trait of . Riggs contends that while some things are epistemically good because of their connection to the epistemic goal of truth, insight—even if it has the effect of leading to truth—is epistemically because of its connection to the distinct epistemic end of understanding. And so for Riggs, while both truth and understanding are epistemic goals, insight is defined (of the latter goal but not the former. This case is a specific example of why pluralism should not be thought of as wedded to the (thesis, and perhaps also, why pluralism would be suspect if it ?(

aiming at knowledge will be the very same purposes we have insofar as we are aiming at truth, or true belief. The idea, then, would be for the veritist to reject an implicit premise of the knowledge objection—the assumption that if our epistemic purposes consist in part in those purposes we have insofar as we are after knowledge, then our epistemic purposes are not those purposes we have insofar as we're after the truth (i.e. maximizing true beliefs, avoiding false beliefs). It was assumed in the argument that these purposes would be distinct enough to not collapse together. Let's call the case for challenging this idea the

E \subset \emptyset (((((((((

1. If we aim at knowledge if and only if we aim at the truth (i.e. true belief), then some purpose P will be a purpose we have insofar as we're after knowledge if and only if P is a purpose we have insofar as we're after truth (i.e. true belief).
2. If some purpose P will be a purpose we have insofar as we're after knowledge if and only if P is a purpose we have insofar as we're after truth (i.e. true belief), then our (are those we have insofar as we're after knowledge if and only if they are also those we have insofar as we're after truth (i.e. true belief).
3. We aim at knowledge if and only if we aim at the truth (i.e. true belief).

4. Therefore, our epistemic purposes are those we have insofar as we're after knowledge if and only if they are also those we have insofar as we're after truth (i.e. true belief).

If the (is sound, then Premises (1) and (2) in the knowledge objection fail to count against veritism because simply ((((((((((would be compatible with our purposes being epistemic only insofar as we're after the truth (the veritist's claim).

The key premise of the (is (3), and (3) is admittedly not without some intuitive support: because knowledge is factive, you can't have as your aim knowledge without also having as your aim truth. And additionally, it seems strange to suppose one could aim to have the truth without aiming at knowledge. After all, if you want the truth about whether or not p, you'll want your belief about p to be as well-justified as it can be; and so you aim for a true belief when you aim for a true belief that is as well justified as it can be. This looks quite a bit like aiming for knowledge.

Despite the appeal of premise (3), the biconditional is actually far too strong, and it's not hard to see on closer inspection how the two aims come apart. First off, consider that even

though true belief is part of what makes up knowledge, true belief does not entail knowledge: they come apart in cases of mere true beliefs. So knowledge and mere true beliefs are distinct. If (3) is correct, then you aim for knowledge only if you are also aiming at a mere true belief. But, to revisit Kvanvig's point:

When asked, "why, from a purely cognitive point of view, do you value knowledge?" we are inclined to answer...that we want to be correct, but not merely by accident, as what happens when one has a merely true belief. (Kvanvig 2009 : 7)

If (3) were correct, then you can take on the aim of knowledge only if willing to admit that you also want to be correct, just by accident 'as what happens when one has a merely true belief' (2009 : 7). Because knowledge excludes luck, you can't both know and be correct by accident. And so (3) is true only if you aim at knowledge only when you also aim at something you have only if you have knowledge. In this respect, (3) is not only false, but in a sense contradictory insofar as it expresses intentional object; consequently, the veritist could rebut the knowledge argument via the collapsing argument only by reasoning through (3), and we can see that this would be a mistake. Here again, the veritist has an initially attractive response to the knowledge argument that ultimately comes unhinged, and irreparably so.

I considered the knowledge objection here in the detail that I did because I wanted to make it clear that the most obvious response to veritism is not one to which it has any clear and ready defense. By tracing out the possible replies, we can see further that veritism lacks not only an obvious defense, but indeed any good defense. It should not seem at this point that veritism is nearly as attractive a position as it appears after just a superficial consideration, or after pointing out the theoretical virtues (i.e. parsimony) that made it look so promising.

Moreover, the fact that the knowledge objection draws out problems with the veritist's proposal should lead us to think that parallel-style objections to veritism that appeal to understanding as an epistemic end will also pose problems. And indeed, Riggs (2004), Brogaard (2005), Elgin (2007) and Kvanvig (2003) among others have offered cases for thinking that understanding is both distinct from truth as an epistemic aim and also one whose epistemic value is not reducible to the value of truth. I'll not consider these objections here, but rather, I think it will be more useful to engage in a different sort of strategy, one that highlights a special sort of problem that itself invites for a new and important set of questions.

5.6 Veritism and Epistemic Responsibility

I want to take a stronger line against the veritist now; in this section, I shall show by example that veritism commits itself to an absurdity. The absurdity itself is designed to reveal something

to us about how the epistemic point of view must be characterized. In that respect, the objective here will not simply be a negative verdict for veritism, but a positive endorsement of the view it will naturally give way to.

I want to begin now with a thought experiment. Imagine that two old men, Gray and White lie upon their deathbeds, praying to God. Watching from above, God gets ready to evaluate Gray and White morally to determine whether they are fit for heaven. Quickly bored, God changes his mind and decides that he wants Gray and White to be evaluated from a purely epistemic, rather than a moral, point of view. He sends an angel, Gabriel, to visit Gray and White, with precisely these instructions. At the bedside, Gabriel is quickly impressed by Gray, who he notices has acquired twice as many truths as White; but upon closer investigation, Gabriel realizes that the majority of Gray's truths were ones he held on the basis of scant evidence, stubbornness and wishful thinking. Gabriel scrawls down his mixed assessment of Gray and next turns to White only to find quite the opposite case. White had, in his life, acquired just half the truths Gray had. 'Not good,' thinks Gabriel, who scratches the number of White's truths in his angel pad. But upon a closer look, Gabriel is stunned: as it turns out, nearly all of White's beliefs were based on well-grounded reasons and on the basis of carefully weighed evidence. It was just a matter of bad fortune that so many of these beliefs turned out false. Gabriel scrawls this information in his angel pad. Seeing Gabriel to have finished, God asks: 'Who was the better epistemic agent? Gray or White?'

It is important that we resist the temptation to smuggle moral considerations into what is here a purely epistemic evaluation; all God cares about here is who the better epistemic agent is. What we're evaluating epistemically then is agency, not belief. Gabriel, suppose, needs a bit of help here; he takes his angel pad to the local veritist, shows the veritist the raw data for both Gray and White and asks who the better epistemic agent was.

Before Gabriel can finish showing the veritist all his notes, and explain again the important fact that God wants to know which (was epistemically better, the veritist interrupts: 'An agent,' says the veritist slowly 'is subject to the same criteria of epistemic evaluation as beliefs, belief-forming processes, and anything else. There is a simple formula that we veritists apply across the board: for any X, X is epistemically valuable to the extent that it maximizes truth and avoids error.' With this, the veritist quickly circles Gray's name and says, 'There's your winner. That agent is epistemically better than the other one.'

Now, I contend here that Gabriel asked the wrong person; the veritist's conclusion is not the right one, and I want to now go about saying why.

When we considered how it is that beliefs aim at truth, we saw that this aim should be understood as characterizing a constitutive fact about belief: part of what it is to be a belief is to be the sort of thing that is correct only when true. Agency, on the other hand, is not like belief in this respect: agency is not (epistemically) correct only when true, and that is so even though truth is something at which agents can aim. If an agent aims at the truth, then, and tries to bring it about: what are the conditions under which her agency is correct? We might be tempted here to say that the epistemic norm that governs agency is truth nonetheless, even though agency itself isn't correct only when true. The idea might be that agency is epistemically correct insofar as it brings about the truth. This is the line the veritist would have us take. But it is misguided. When truth makes beliefs correct, it is because beliefs fit the world. When does truth make agency correct? When agency 'fits the world'? Neither the intentional aims of agents nor agency itself is something that aspires to fit the world. The direction of fit is the other way around. Perhaps then agency is epistemically correct when the world fits the agent's intentional aims? Even if one is guided by the goal of truth, truth cannot be an intentional aim. As Lynch (2008) says:

One does not simply will oneself to believe the truth. Rather, we pursue truth indirectly, by pursuing those beliefs backed by reasons and supported by the evidence. (Lynch 2008: 4)

It stands to reason then that agency is epistemically correct just when the world fits the agent's intentional aims of having beliefs backed by reasons, supported by evidence, etc. The world fits these aims, and thus agency will be correct, just when the agent's succeeds in her aim of believing responsibly.

The picture this gives us then is one according to which beliefs aim at truth and agents aim at responsible belief. The fundamental aim of belief is truth, the fundamental aim of epistemic agency is responsible belief. On this proposal, we can make much better sense of the case of Gray and White than the veritist could; our view says that if it's (that are being evaluated from the epistemic point of view, then Gray's are more epistemically valuable than White's. This seemed to be the right result, as Gray had far more true beliefs. If it's agency that is being evaluated, however, our view clearly rules that White, rather than Gray, is the better epistemic agent. The fundamental aim of epistemic agency is responsible belief, and White had far more responsible beliefs than Gray. The present view is, thus, a variety of (, although not in the standard sense: unlike pluralists such as Riggs, I am not allowing for manifold sources of epistemic value—instead, there will be just one epistemic value to govern belief, and another to govern agency.

5.7 Knowledge and the Epistemic Point of View

But before moving on, there is a final conclusion that we are in a position to draw—a conclusion about what sort of thing knowledge is. Here it will be helpful to retrace our steps: our overarching project has been to develop an anti-luck theory of knowledge out of our value platitude. Up to now, our aim has been a modest one—simply to find out what exactly the Meno Requirement demands of us, and that has been a task in itself. What we found in Chapter 4 was that the Meno Requirement consists in two distinct criteria, one pertaining to knowledge and true belief qua types and the other pertaining to knowledge and true belief qua tokens. But because we saw that the Meno Requirement is one that we can't purport to address without a theory of epistemic axiology, we needed to specify one that would work. We concluded in this chapter that veritism is inadequate, and that we should understand the epistemic point of view as broadly one on which both the good by way of belief and by way of agency matter. We know then that a theory of knowledge that preserves the two criteria of the Meno Requirement is going to have to do so in light of the fact that both the good by way of belief and agency are fundamental epistemic values. A precedent for this sort of thinking is found in Sosa's (2007) recent work on epistemic normativity, whereby Sosa supposes a response to the value problem for knowledge turns on the thought that knowledge is connected to distinct though equally fundamental epistemic values. Here's Sosa:

One part at least of the solution to the value problem lies in a point central to virtue epistemology: namely, that the value of apt belief is no less epistemically fundamental than that of true belief. For this imports a way in which epistemic virtues enter in the attainment of fundamental value, not just instrumentally. (Sosa 2007: 87)

Sosa here recognizes that the epistemic value of virtuously formed belief should play a role in accounting for the value of knowledge not in strumentally as a means to truth, but as a fundamental epistemic value in its own right. Does this sort of dual account of fundamental epistemic value tell us anything about what a theory of knowledge will look like? I think it does. By learning all that is entailed in satisfying the Meno Requirement, and determining at least roughly what the value sources are (good belief and good agency) within which the Meno Requirement is framed, we're in a position to see what the Meno Requirement implies that knowledge be . For one thing, knowledge can't be:

K1: A true belief that is not held responsibly.

(2) To develop this template into a defensible account of knowledge. Because knowledge excludes luck, the core task in doing so will be to show how the connection between true belief and the cognitive ability it must arise from will suffice to block the relevant sorts of epistemic luck.

The consequence of satisfying these objectives, then, will be that we'll have shown how the connection between true belief and cognitive ability essential to knowing is one that both seals knowledge from undermining luck and accounts for the distinctive value we take knowledge to have. This will be, ultimately then, a theory of knowledge that is not simply developed independently from the luck and value problems and only later tested against them; it will be a theory that builds a satisfaction of the conditions arising from these problems into an analysis of what knowledge is. (