

Veritism Unswamped

KURT SYLVAN

University of Southampton

k.l.sylvan@soton.ac.uk

According to Veritism, true belief is the sole fundamental epistemic value. Epistemologists often take Veritism to entail that all other epistemic items can only have value by standing in certain instrumental relations—viz., by tending to produce a high ratio of true to false beliefs or by being products of sources with this tendency. Yet many value theorists outside epistemology deny that all derivative value is grounded in instrumental relations to fundamental value ('Instrumentalism'). Veritists, I believe, can and should follow suit. After setting the stage in §1, I explain in §2 why Veritism should not take an Instrumentalist form. Instrumentalist Veritism faces a generalized version of the swamping problem. But this problem undermines Instrumentalism, not Veritism: granting Instrumentalism, similar problems arise for *any* economical epistemic axiology. I show in §3 how Veritism can take a less narrow form and solve the swamping problem. After answering some objections in §4, I consider what some would regard a less radical alternative solution and argue that it either fails or collapses into mine. I close in §5 by taking stock and re-evaluating the overall prospects for Veritism, suggesting that it is a highly promising epistemic axiology when divorced from Instrumentalism.

1 Introduction

'[T]he proper appreciation of a beautiful object is a good thing,' wrote G. E. Moore (1903, Ch. VI, §114). Yet while beauty and its proper appreciation are both good, it is implausible that they are equally *fundamental* goods. Appreciating beauty is good only because beauty is good. After all, it is only good to appreciate what merits appreciation. Appreciating trash is no good. So, although it is valuable, the appreciation of beauty seems to have a *derivative* kind of value.

This is an example of a more general fact. In any evaluative domain, some values are more fundamental than others, in the sense that the value of everything else in the domain is explained by their value. So, for any domain, we can ask the *Fundamentality Question*:

(FQ) Which value or set of values is the most fundamental in the domain?

Ethicists have long addressed FQ as it arises in the practical domain. Epistemologists have recently taken interest in FQ as it arises in the epistemic domain. There are, after all, many epistemic values: accuracy, rationality, justification, coherence, knowledge, etc. But it is doubtful that they are equally fundamental. We admire some of them from the epistemic point of view because we admire others from the epistemic point of view.

So, which is fundamental? Many have found it attractive to think that truth is at the bottom of it all. For many items we value from the epistemic point of view, it is plausible that we value them *because* we value accurate belief. Inspired by this idea, one might favor:

(*Veritism*) Accurate belief is the sole fundamental epistemic value.

Many epistemologists have accepted Veritism, though it finds considerable opposition in recent literature. I also accept the view. But I think opponents and proponents alike understand it in an unjustifiably narrow way. Veritism is defensible if and only if it is understood

less narrowly. The aim of this paper is to explain and defend this claim.

The Narrow Assumption: Instrumentalism about Derivative Value. Many epistemologists assume that there is only one kind of way in which we can explain one epistemic value in terms of a more fundamental epistemic value. The explanation, they assume, must proceed by invoking *instrumental relations*, so that for any X, X is derivatively epistemically good only by (i) tending to produce fundamental epistemic goods or (ii) being the product of something with feature (i).¹ Call the kind of value something has just in virtue of (i) *production value*, and the kind something has just in virtue of (ii) *mere product value*. More officially, then, many assume:

(*Instrumentalism about Derivative Epistemic Value*) All derivative epistemic value is either production value or mere product value.

As a result, many assume that Veritism is trivially equivalent to:

(*Instrumentalist Veritism*) Accurate belief is the sole non-instrumental epistemic value.

None of this is trivial, however. Many value theorists have rejected Instrumentalism, including some who are consequentialists about rightness. Accordingly, it is only natural to wonder whether Veritism can take a less narrow form.

What could derivative value be, if not instrumentally grounded value? The example with which I opened provides one illustration. Appreciating beauty is good *because* beauty is good. But this ‘because’ signals no instrumental relation. Appreciating beauty does not reliably cause more beauty to exist, and needn’t be the product of anything beauty-conducive. Luckily, we are not forced to appeal to instrumental relations. On a more natural model, the instance of appreciation derives value because (a) its intentional object is good and (b) it is a proper way to value that good object. The result is appealing: appreciating beauty seems parasitically but also non-instrumentally good.

Here is another way to understand what is going on. For X’s value to be explained by Y’s, X must stand in some relation to Y. Call such relations *value derivation relations*. Instrumental relations comprise one species of value derivation relations. But not the only species. In the case of beauty and its appreciation, the relation that mediates the derivation is not any instrumental relation but rather the relation of *being a way to value*.

Hurka’s Principle. Ethicists have constructed axiologies that honor this idea. Thomas Hurka, for example, defends a recursive account on which there are *basic* non-instrumental values such as beauty and pleasure, and a principle for grounding *derivative* non-instrumental values in the basic ones:

(*Hurka’s Principle*) When V is a non-instrumental value, proper ways of valuing V have some derivative non-instrumental value.²

¹One might reserve ‘instrumental epistemic value’ for (i). But I use it to capture the kind of derivative value that something has *just by standing in some instrumental relation*. The *product of* relation is just as much an instrumental relation as the *caused by* relation is a causal relation. Since many epistemologists allow that (ii) can ground derivative epistemic value, it is only fair to define Instrumentalism this way.

²See Hurka (2001). Hurka uses a less general version of this principle which mentions *love* rather than *ways of valuing*. But that principle is, I take it, plausible because love is a proper way of valuing the good.

Is the principle trivial because ‘proper’ and ‘valuable’ are synonyms? No: propriety and value are distinct normative categories. The principle also says proper ways of valuing are *derivatively* valuable, not just valuable.

Given this principle, Hurka suggests that consequentialists about rightness can embrace the non-instrumental value of virtue while also respecting the thought that virtue has only a parasitic kind of value. They can do so by identifying virtues with ways of valuing more fundamental values and by making use of Hurka's Principle. Hurka's model is not the only one in value theory that rejects Instrumentalism. Other axiologists reject it.³ I highlight Hurka's view because it inspires my own.

Before proceeding, I want to make some clarifications about how I understand Hurka's Principle. Firstly, I construe 'valuing' broadly, to include any way of being positively oriented toward something in attitude, act, or disposition. I do not have in mind the narrow sense that contrasts with first-order desiring.⁴ As examples of valuing in this broad sense, I would include everything picked out by Nozick's 'V verbs' and more, which include:

maintaining, saving from destruction, prizing, contemplating, [...] car[ing] about, accept[ing], support[ing], affirm[ing], encourag[ing], protect[ing], guard[ing], prais[ing], seek[ing], embrac[ing], serv[ing], be[ing] drawn toward, be[ing] attracted by, aspir[ing] toward, striv[ing] to realize, foster[ing], express[ing], nurtur[ing], delight[ing] in, respect[ing], be[ing] inspired by, tak[ing] joy in, resonat[ing] with, be[ing] loyal to, be[ing] dedicated to, and celebrat[ing]. (Nozick 1984, p. 429)

This broad understanding fits both with the way proponents of Hurka's Principle intend it and with the core intuitions that support it. While Hurka uses the narrow-sounding 'love' to express his principle, he is clear that he intends this word to be understood unusually broadly, so that mere desiring counts as loving.⁵ Similarly, Adams talks about 'being for' the good, and counts as ways of being for the good many of the phenomena on Nozick's list (Adams 2006, pp. 15–16). It is sensible, I believe, to understand the principle this broadly. While it might be *better* to value the good in a strong sense that implies higher-order endorsement, even mere first-order aiming at the good is somewhat non-instrumentally good. Consider Huck Finn, who is motivated by a desire to help Jim that he rejects as wrong at the higher order.⁶

Hurka's Principle is hence very general. I think it is natural to generalize it in further ways that will matter in what follows. To see another plausible generalization, note that it is not only ways of valuing (e.g., love) but also actions and attitudes that *manifest* them (e.g., loving treatment) that have derivative non-instrumental value.

By way of illustration, suppose Alice performs a beneficent act that manifests her valuing of beneficence, while Beatrice performs the same kind of act as a PR stunt. Alice's act has greater worth than Beatrice's, and this extra worth is not merely instrumental. But the fact that Alice's act has greater worth is not just a brute fact. It has a two-stage explanation: (i) Alice's act has greater worth by manifesting real valuing of beneficence, and (ii) this valuing, in turn, is good because it is fittingly directed at something good. So, we should extend the principle to say:

When V is a non-instrumental value, fitting ways of valuing V *and their manifestations* have some derivative non-instrumental value.

³See Anderson (1993), Herman (1993), Kagan (1999), Korsgaard (1983), Moore (1903), Parfit (2011), Rabinowicz and Ronnow-Rasmussen (1999), Rashdall (1907), Scanlon (1998), and Zimmerman (2001).

⁴The restricted sense of 'valuing', which is admittedly often salient when we use the term in ordinary language, came into prominence in action theory after Watson (1975), and played an important role in discussions of subjectivism about value by Lewis (1989), Harman (1993), and Smith (1994).

⁵Note Hurka's gloss: 'To 'love' x is to be positively oriented toward x in one's desires, actions, or feelings, or, more generally, in one's attitudes' (2003, p. 14).

⁶Cf. Arpaly (2003).

One further extension will prove important. Like Sosa (2007), I will only be assuming that accuracy, knowledge, rationality, etc., have value *from the epistemic point of view*. So, I am only interested in defending a version of Veritism on which accuracy is the sole fundamental value from the epistemic point of view. It is compatible with this view that accuracy might have no non-instrumental value from any *non-epistemic* point of view. Thus, I need the principle that if V-ing is a fitting way of valuing a non-instrumental value *from the point of view of a domain D*, V-ing has derivative non-instrumental value *in D*.

This extended principle is no less plausible than the original. Even if beauty isn't non-instrumentally valuable *simpliciter*, it is non-instrumentally valuable *from the aesthetic point of view*. So is the appreciation of beauty: this is also a good thing from the aesthetic point of view, since it is an aesthetically fitting response.⁷ But the latter has a derivative status relative to the former. So, we should revise the principle in one further way:

(*The Extended Hurka Principle*) When V is a non-instrumental value from the point of view of domain D, fitting ways of valuing V in D and their manifestations have some derivative non-instrumental value in D.

A Caption of the View and the Plan. My view appeals to the Extended Hurka Principle to explain why rational belief, justified belief, and knowledge have an accuracy-oriented kind of epistemic worth. Such beliefs are epistemically valuable because they manifest certain ways to value accuracy in thought. The view is compatible with many substantive pictures of the nature of rationality, justification, knowledge, and other epistemic desiderata. I see the best pictures as simply disagreeing about what it takes to value accuracy in thought. My view is that any way of valuing accuracy in believing P involves a disposition to hold a belief in P only if it is likely accurate relative to the epistemic reasons. There are several ways to value accuracy in thought because there are several kinds of likelihood and epistemic reasons, corresponding to different epistemic desiderata. These ways and the corresponding desiderata qualify as derivatively epistemically good thanks to the Extended Hurka Principle.

With this capsule statement in mind, here is the plan. In §2, I explain why Veritism should not be identified with Instrumentalist Veritism. Instrumentalist Veritism faces a generalized version of the swamping problem. But this problem undermines Instrumentalism, not Veritism. Granting Instrumentalism, similar problems arise for *any* economical epistemic axiology. I show in §3 how Veritism can take a less narrow form and solve the swamping problem. After answering some objections in §4, I consider what some would regard as a successful and less radical alternative solution—the solution offered by virtue epistemologists—and argue that it either fails or collapses into mine. I close in §5 by taking stock and re-evaluating the overall prospects for Veritism. While it is beyond the scope of this paper to establish that Veritism is hands-down superior to every alternative epistemic axiology, I suggest that Veritism renews its promise when divorced from Instrumentalism.

Disclaimers. Before moving on, let me make some remarks about what I'm *not* doing. Denying Instrumentalism is compatible with affirming that some derivative epistemic values are purely instrumental. Reliable types of processes have significant instrumental epistemic value, and reliabilists are right to place value on them. I don't deny these claims.

Denying Instrumentalism is also compatible with the idea that some items might have both instrumental epistemic value and a different kind of derivative epistemic value. Being

⁷Of course, it sounds weird to say that the appreciation of beauty is aesthetically good. But the problem is terminological: since appreciation is an aesthetically fitting response to beauty, it is certainly a good thing from the aesthetic point of view. This is why I use the 'point of view' talk, which is more general.

instrumental to accuracy, I agree, is a *necessary* condition for some epistemic virtues. I only deny that Instrumentalism can explain all the facts about derivative epistemic value.

Furthermore, I am not arguing in this paper against reliabilism. Reliabilism is a thesis about the nature of justified belief or—in some permutations—knowledge. Instrumentalism is a thesis about how derivative epistemic value is grounded in fundamental epistemic value. Reliabilism does not entail this thesis. Indeed, reliabilism does not by itself entail any claims about the worth that any beliefs possess from the epistemic point of view. One would need to appeal to auxiliary assumptions about epistemic worth and the relationship between derivative and fundamental epistemic worth to draw any such conclusions.

Many have worried that reliabilists can *only* appeal to an Instrumentalist model of derivative epistemic value. But I think this is wrong: there are sophisticated forms of reliabilism that are not forced to accept a purely Instrumentalist model. Admittedly, I will discuss a simple form of reliabilism when I introduce the generalized swamping problem for Instrumentalism in the next section. But this is only as a nod to the literature.

Finally, my focus is on epistemic *value*. Here I will not defend a non-consequentialist view about epistemic *rightness*. Indeed, Hurka's view was intended to provide a way for consequentialists about ethical rightness to consistently uphold the non-instrumental value of virtue. By accepting a richer axiology, consequentialists can uphold a less revisionary account of our duties. This idea is old, found in ideal consequentialists of the early 20th century like Hastings Rashdall and G. E. Moore. While I do reject consequentialism about epistemic rightness, this paper is not an argument against it.⁸

2 Instrumentalism and the Swamping Problem

Why should Veritists reject Instrumentalism? A large reason is that this is the best way to solve a generalized version of the swamping problem. To bring this out, I will rehearse the original swamping problem and explain why we should still take it seriously. I will then explain why it undermines Instrumentalist Veritism. But I'll argue that the Instrumentalist half is the culprit: *any* modest epistemic axiology that embraces Instrumentalism faces a relative of the swamping problem, and Veritists can easily avoid it by rejecting Instrumentalism.

2.1 The Old Problem (and Why It Remains Important)

Originally, the swamping problem was presented as a problem for a simple kind of reliabilism. The following thoughts prompted the problem. The epistemic value of a reliable belief-forming process *per se* is just instrumental epistemic value relative to the goal of producing true beliefs. While reliable belief-forming processes have great epistemic value, one also wants to evaluate their products. And unfortunately, the following claim is plausible:

(A) A belief's having been produced by a reliable belief-forming process does not *as such* make that belief epistemically better if that belief is already true.

(A) is made plausible by an analogy from Zagzebski (1999). The mere fact that some good coffee was produced by a reliable coffeemaker does not make that coffee better. But a reliably produced true belief as such is analogous to a reliably produced cup of good coffee. Thus:

(B) If knowledge = true belief produced by a reliable belief-forming process, then knowledge is not as such epistemically better than true belief.

⁸See Berker (2013) for an argument against consequentialism about epistemic rightness.

But knowledge is as such epistemically better than true belief. Hence the problem.

Some might try to resist this argument by observing that we do sometimes place greater value on products of reliable sources.⁹ For example, the market value of a good watch produced by Rolex would be far higher than that of a qualitatively identical watch produced by Casio. Doesn't this undermine the swamping argument?

No. Intuition pumps like this do support the conclusion that some products of reliable sources are better than intrinsically similar products of unreliable sources. But this conclusion does not undermine the reasoning behind the swamping problem. Here is why. The key thought behind the swamping argument is this:

(*) Being produced by a reliable producer of good Fs cannot *as such* make a good F better.

The 'as such' matters. It is consistent with (*) that other properties can *contingently accompany* a good F's reliable ancestry and add value to that F. The Rolex/Casio example is an illustration: Rolexes are status symbols and have a higher resale value. That is why the Rolex is preferable. If it lacked these properties, preferring it would be irrational.

Proponents of the swamping argument will agree that reliably produced good Fs are sometimes better than unreliably produced good Fs. Zagzebski (1996) agrees that reliability is a necessary condition for complete epistemic virtue.¹⁰ Thus, she agrees that some reliably produced true beliefs are better than mere true beliefs: some reliably produced true beliefs will satisfy the *other* necessary conditions for manifesting complete epistemic virtue, after all! What Zagzebski and others deny is that reliable ancestry *as such* is sufficient to explain why knowledge *as such* is epistemically better.

That was the problem all along. So, the conclusion (i.e., (B)) from the original literature stands. It is insufficient to observe that a reliably produced F is sometimes better than an unreliably produced but intrinsically similar F. Zagzebski and others will agree but deny that the explanation proceeds *via* the bare fact of reliable production. That was the point of the coffee analogy. The analogy forces us to ask what else could make the difference, since reliable ancestry *as such* seems insufficient.

This isn't to deny that reliable processes have great epistemic value. The problem has always concerned their products. As Ward Jones put it:

When we ask about the value of knowledge over mere true belief, we are asking not about belief-forming methods but about their products. We have clear reason to care about reliable methods of belief-formation, and the reliabilist is right to emphasize them. But it is unclear, on reliabilism, why we should divide the desired products of belief-formation, true beliefs, into those which were brought about by reliable methods and those that were not. The epistemic instrumentalist gives us no way of evaluatively dividing true beliefs into those which have been brought about by justifying methods versus those which have not (Jones 1997, p. 425).¹¹

This challenge remains a serious one.

⁹Cf. Goldman and Olsson (2009) and Goldman (2012).

¹⁰See Zagzebski (1996, pp. 165–194).

¹¹Interestingly, Jones traces the swamping problem to truth-instrumentalism and notes that it is a problem for internalists as well as reliabilists. But he didn't consider how one might reject instrumentalism while preserving Veritism, or emphasize other forms of value derivation. Zagzebski similarly emphasizes the generality of the problem and traces it (in effect) to truth-instrumentalism, but again doesn't think Veritism can be preserved (Zagzebski 2003, p.16); indeed, she opposes Veritism there and monism generally in Zagzebski (2004). Zagzebski and Jones deserve credit for their early appreciation of the problem's generality. But they didn't observe the core axiological point that one can remain monistic on fundamental value if one is pluralist about value derivation. It is surprising, since Zagzebski acknowledges that given that 'if something is valuable it is also valuable to appreciate or love it, then love of true belief has value because true belief has value' (Zagzebski 2003, p.18).

2.2 The Deeper Problem

The swamping problem runs deeper, as others have recognized.¹² It is easy to see that the problem is not just a problem about knowledge or just a problem for reliabilists.

To see the first point, note that reliabilists identify a belief's being justified with its being produced by a reliable type of belief-forming process. Assume they are right for the sake of argument. We can use the same coffee analogy to argue that *justification* cannot as such add epistemic value to true belief. That is bad: justification as such does add epistemic value to true belief! This is not a restatement of the problem about knowledge: due to the Gettier problem, no reliabilist will equate knowledge with justified true belief.

There is an even more general structure that makes the problem not just of limited interest to reliabilists. On any view on which

being justified : true belief :: being produced by a good coffeemaker : good coffee

there is a worry that justification cannot as such add epistemic value to true belief. So, there is a worry for any view on which the epistemic value of justified belief consists in its being the *mere product* of some type of source that is only instrumentally good relative to true belief. Reliabilists are hardly the only epistemologists who accept this assumption.

These observations lead to a general argument against Instrumentalist Veritism.¹³ Instrumentalist Veritism entails that:

- (I) Items other than true belief are epistemically good iff they (i) tend to produce true (and not false) beliefs or (ii) are products of a source with property (i).

And the following is a natural generalization of the points about the coffeemaker case:

- (II: *Swamping Premise*) If X has its source in something that is only instrumentally good relative to property F and X already exemplifies F, the mere fact that X has that source cannot as such make X better.

But according to Instrumentalist Veritism, justifying sources only have instrumental epistemic value relative to true belief. So (I) and (II) will entail:

(*Bad*) A true belief's being justified cannot as such make it epistemically better.

Some would take this to undermine Veritism.¹⁴ But this is only because they assume Instrumentalism. Veritists can reject Instrumentalism and view the generalized swamping problem as an argument against it. To support this tactic, I will show that the Swamping Premise does not extend to other species of derivative value. This will show that the problem is a local one for Instrumentalists. To drive the point home, I will show that there are similar problems for all modest alternatives to Veritism that cling to Instrumentalism.

2.3 Why Other Forms of Derivative Value Help

Let us first understand why rejecting Instrumentalism can help Veritists. Doing so can help them, I suggest, because other forms of derivative value are not subject to swamping by the values on which they are parasitic. To see this, consider:

¹²Cf. Pritchard (2010, 2011).

¹³As Duncan Pritchard (2010, 2011) in effect observed.

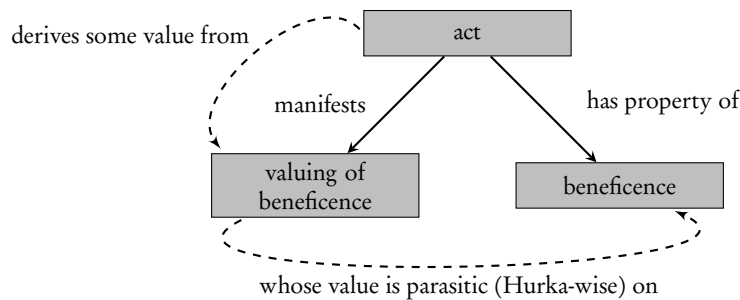
¹⁴Cf. Pritchard (2011).

(Stronger Swamping Premise) For no type of derivative value and no sense of ‘has its source in’ is it true that:

if X has its source in something that only has derivative value relative to property F but X already has F, X’s having that source *as such* makes X better.

This is false. Suppose Alice performs an act of beneficence because she values beneficence, while Beatrice performs the same kind of action as a PR stunt. Plausibly, Alice’s action is better than Beatrice’s because it manifests (and in *this* sense has its source in) real concern for beneficence. Yet just as appreciating beauty is good *because* beauty is good, so valuing beneficence is good *because* beneficence is good.

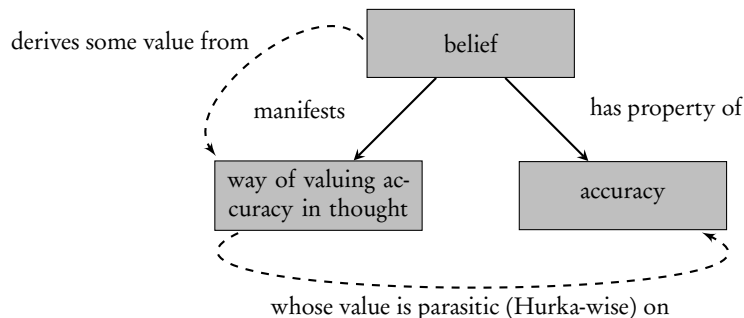
Here an act derives value by manifesting something only derivatively good relative to another property that the act exemplifies. The source manifested is (1) Alice’s valuing of beneficence, and the property exemplified is (2) beneficence. Yet it is plausible that the sheer fact that Alice’s act has its source in (1) makes that act worthier *per se*. In a picture:



This would be impossible if the Stronger Swamping Premise were true. So it is false.

If the foregoing points are right, Veritism generates no problem: only its conjunction with Instrumentalism does. If Veritists adopt a different model of value derivation, they can avoid swamping just as it was avoided in the case of Alice’s beneficent action.

On the view I develop in §3, rational belief and knowledge gain epistemic worth by manifesting different ways of valuing accuracy in thought. On this view, the case of knowing parallels the case of Alice’s beneficent action. Knowledge requires a belief whose accuracy manifests a disposition to hold beliefs only if there is sufficient objective and subjective evidence that they are true. This disposition *just is* a way to value accuracy in thought. Ways of valuing accuracy are epistemically good because accuracy is epistemically good. But this ‘because’ is not (purely) instrumental. Paralleling Alice’s case:



I will explain the ways of valuing accuracy in thought in §3. The tactic is what matters now. Veritists can claim that true beliefs can gain epistemic worth by manifesting ways of valuing accuracy, just as beneficent actions can gain moral worth by manifesting an agent's valuing of beneficence. Both cases undermine the Stronger Swamping Premise.

So Veritists can avoid the swamping problem if they appeal to forms of derivative value beyond the instrumental. But they must avoid the problem in this way. For the restricted Swamping Premise *is* true. Instrumentalist Veritism *is* false.

2.4 A General Problem for Instrumentalists

We can strengthen this advice by seeing that if Instrumentalism were true, other economical epistemic axiologies would face relatives of the swamping problem.

Observe that there is another side to Zagzebski's analogy. We do not only think that being produced by a reliable coffeemaker *per se* cannot improve good cups of coffee. We also think that such ancestry cannot improve bad cups of coffee. If you drink vile coffee and cringe, it is not comforting to be told: 'Hey, at least it was produced by a reliable coffeemaker.' These facts support a more striking sibling of the restricted Swamping Premise:

(*Dud Principle*) If X was produced by a source that is only good because it produces good Fs, that fact as such can't make X better if X is otherwise a bad F.

Carter and Jarvis (2012) took this to undermine the intuitions behind the swamping problem. For they thought that the Dud Principle would imply, crazily, that "non-factive epistemic properties—most saliently justification—are never epistemically valuable" (Carter and Jarvis 2012, p.690). But this is too fast: that conclusion follows only if we grant Instrumentalism. Like the original Swamping Premise, the Dud Principle concerns mere product value. The crazy conclusion would not follow from the Dud Principle if non-factive epistemic properties had a different kind of derivative epistemic value. Only if the Dud Principle extends to other kinds of derivative value is Carter and Jarvis's conclusion fair.

But the Dud Principle does not generalize, for the same reasons why the Swamping Premise did not generalize. Consider someone trying to perform a beneficent act as a manifestation of her valuing of beneficence, failing only due to bad luck. Her efforts remain better than the failed efforts of someone merely looking for a PR boost.

What the Dud Principle really suggests is that Instrumentalism is an incomplete model of derivative epistemic value. After all, it is not as if Carter and Jarvis can convince us that vile coffee is better if it comes from an otherwise reliable coffee machine. Like the original Swamping Premise, the Dud Principle captures a fact about mere products of instrumentally good sources. But it is then easy to see that any modest axiology that embraces Instrumentalism will face a relative of the swamping problem.

Suppose knowledge is one's fundamental epistemic good. Given Instrumentalism, how can one explain the epistemic value of justified beliefs? One must claim that such beliefs are good by being products of knowledge-conducive processes. But the Dud Principle then makes it mysterious why justified false beliefs are epistemically good. From an Instrumentalist point of view, they are duds just like bad cups from otherwise reliable coffeemakers.

Expanding the stock of fundamental epistemic values doesn't really help. Even if one adds rational belief, justified belief, knowledge, understanding, and true belief to the list, there remain epistemic values that (a) no one can reasonably take to be fundamental, but (b) admit of no Instrumentalist explanation. Consider the epistemic value of *trying one's best to form beliefs accurately*. This is a paradigmatically derivative value: trying to do something

good is admirable because it is intentionally directed at something good. But if we accept Instrumentalism, it is hard to explain why it is derivatively epistemically good. Merely trying to form one's beliefs accurately is not reliably instrumental to accuracy, knowledge, justification, etc. Yet there remains something admirable in one's best efforts.

Any modest epistemic axiology that endorses Instrumentalism will face a relative of the swamping problem. Some paradigmatically derivative epistemic values admit of no general Instrumentalist explanation. So, everyone, not just Veritists, should reject Instrumentalism.

3 Non-Instrumentalist Veritism

Of course, work remains for Veritists. They must provide a more specific view that takes advantage of our observations about forms of value derivation beyond the instrumental, and explain how this view can capture central intuitions about epistemic value.

My view will secure these desiderata. After noting that there are more ways to value something than by promoting it, I will suggest that central epistemic values can be viewed as manifesting different ways of valuing accuracy, and as deriving accuracy-oriented epistemic value thanks to the Extended Hurka Principle.

3.1 Other Ways of Valuing in General

Let's start with a more general fact. It is a truism that values are items that it is proper to value. Given the truism, one should ask for any fundamental value V in any domain:

Which ways of valuing V are the proper ways in the domain?

Instrumentalists might accept a narrow answer:

(The Teleological Answer) For any fundamental value V , the only basically proper way to value V is to instrumentally promote V .¹⁵

But the Teleological Answer is implausible. There are many ways to value: dedication, loyalty, respect, veneration, love, support, and the many other items on Nozick's list. This answer regards only one as basically fitting to fundamental value. This is hardly a default view. For some values, promotion is not the basically proper response. If friendship were fundamentally 'to be promoted', we could properly spend less time caring about the friends we have and more time amassing friends or causing others to have more friends. But we can properly value friendship without taking ourselves to have reasons to produce more instances of friendship.¹⁶ Fans of the Teleological Answer could multiply fundamental values in reply, and claim that loyalty, commitment, respect, dedication, etc., are all fundamental values. But this pluralism is implausible. Besides violating canons of parsimony, it fails to explain asymmetries. Dedication to one's friends matters *because* friendship matters.

These points highlight an insight needed for a version of Veritism that rejects Instrumentalism. The Veritist should capitalize on this insight and suggest that there are more ways to value accuracy in thought than by producing a high ratio of true to false beliefs by any means, including means that disrespect accuracy.

¹⁵I say 'basically proper' because the instrumentalist can obviously allow that other ways of valuing V are non-basically proper in virtue of helping to instrumentally promote V .

¹⁶Cf. Scanlon (1998: 89).

3.2 Ways to Place Value on Accuracy in Thought

How can we implement the insight? What are the ways to value accuracy in thought? I understand the ways that are suited to play a role in traditional epistemological projects as different positive relations to the following norm of accuracy:¹⁷

(AN) It is correct to believe P iff P is true.

AN is not a directly belief-guiding norm. But this norm can guide us indirectly.¹⁸ In being so guided, we value accuracy in our thinking. Coherent beliefs, rational beliefs, justified beliefs, and knowledge can be viewed as epistemically good from a truth-oriented point of view by manifesting different ways to value accuracy in thought.

I will identify four positive relations to the norm of accuracy we can use to explain the epistemic value of coherent, rational, justified, and knowledgeable belief from an accuracy-oriented perspective. Before doing so, I will draw attention to a larger spectrum of ways of being properly oriented toward values within which these fall. Since our ordinary talk about ways of valuing is rather fuzzy, I will have to regiment the language a bit and stipulate narrower, technical uses for some ordinary terms. But it will be clear, I hope, that there are real joints on a pre-theoretically recognizable spectrum that these uses carve out. To further forestall the worry that the discussion is merely stipulative, I will stress the connections that my notions bear to ordinary evaluations of blameworthiness and praiseworthiness.

3.2.1 Proper Valuing in Disposition and Action: A Spectrum

I begin with the obvious: our dispositions and acts can manifest a more or less proper valuing of some value or norm, and they will be correspondingly more or less admirable. Taking the value of equality as an example, let's start at the low end of the spectrum.

Consider Alice, who is disposed to treat people in the ways that she *believes* are equal. There is something *pro tanto* admirable about Alice's disposition. Its *pro tanto* admirableness plausibly owes to the fact that it constitutes a minimal way of valuing equality. Consistently with this fact, Alice might fail to value equality in a fully proper way, and fail to be fully admirable. Suppose her beliefs about what would constitute equal treatment are ones that she would see to be mistaken—and would modify and heed—if she reflected longer on facts she is in a position to know. Her valuing seems less than fully proper. Nevertheless, provided that her ignorance is not deliberate, she manifests *some* consideration for equality, and there is *some* good in this consideration.¹⁹ While it would be better if she took more care to fit her beliefs about equality to her evidence, it would also be worse in one way if she didn't manifested no consideration for equality at all. She is at least well-intentioned, and while this might not count for much, it is *pro tanto* better to be well-intentioned than not to be.

¹⁷I make no claims about ways of valuing accuracy that are not suited to play a role in these projects. Here I distance myself from responsibilists less conservative than Zagzebski. I also distance myself from any standard kind of responsibilism, since responsibilists typically stress intellectual character traits. My ways of valuing are not character traits. Here again I follow Hurka, who divorces virtue from character in his (2006).

¹⁸Cf. Wedgwood (2002, 2013).

¹⁹I wouldn't make claims of the same strength as the one that follows about cases of motivated ignorance, such as the cases Moody-Adams (1994) discusses. These cases involve people who manifest more disregard for value than regard for value. Exactly how spinoffs of Alice's case involving motivated ignorance bear on Hurka's Principle is a delicate question. Consistently with Hurka's Principle, we could say that these spinoffs would involve merely a low level of *pro tanto* intrinsic value that is outweighed by a greater level of *pro tanto* intrinsic disvalue generated by the disregard manifested.

Now compare another agent, Beatrice, who has a disposition not merely to treat people in the ways that she believes are equal, but also to treat people in the ways her evidence suggests comply with the reasons of equality. Beatrice's valuing of equality is more admirable than Alice's. Yet even her disposition may fall short of fully proper valuing. The point is familiar from recent literature on moral worth.²⁰ Beatrice's evidence might be problematically limited, and not necessarily in a way that reflects a procedural mistake in evidence gathering. Suppose, for example, she lives in a benighted age and community in which it is widely believed, often claimed by her epistemic superiors, and supposedly substantiated by scientific evidence that human beings in a certain group lack some essential feature of personhood. No matter how hard Beatrice reflects on the information available to her, she will come to believe these individuals are not persons. As a result, Beatrice is not disposed to treat members in this group with the same level of consideration that she treats other human beings, though she is disposed to treat all individuals that her evidence suggests are persons equally. While Beatrice is perhaps less flawed than Alice, she is hardly flawless. While she is more than merely well-meaning, she still does not value equality in a fully proper way.

Now consider a third agent, Candace, who has a yet stronger disposition. Candace is disposed to treat individuals in the ways that are actually supported by reasons of equality, and for these reasons. She is sensitive not merely to apparent reasons of equality, but to the objective reasons that underwrite the value of equality.²¹ Candace manifests excellence in being for the value of equality, as Adams (2006) would say. Admittedly, her disposition could be more or less steadfast. Even if she is largely sensitive to objective reasons of equality, Candace might sometimes fail to conform to them. Indeed, dispositions are relative, as Sosa (2015) says, to suitable situations and shape, and even in appropriate conditions a subject might exhibit more or less excellence in responding to the equality-relevant reasons. So Candace might fail to do what these reasons support in a particular case. In this respect we might contrast Candace with Doris, who infallibly conforms to these reasons.

In fact, even if a person manifests her sensitivity to a value and successfully does what that value favors doing, her action might still fall short as a manifestation of her valuing because its *success* is not explained by her sensitivity. To take a different kind of example, suppose the reasons of friendship favor Doris's getting her friend Jim a gift that he will enjoy on his birthday. And suppose Doris learns from a reliable source that Jim likes the films of Nicholas Roeg, and so gets him a copy of *The Man Who Fell to Earth*. As it turns out, the reliable source is wrong: her friend has only seen one Roeg film and didn't particularly like it. But in the last couple of months while Doris was travelling and out of contact with Jim, he has become obsessed with the mid-1970s David Bowie. Accordingly, there is a good reason to get John this gift. Here Doris manifests her general sensitivity to reasons of friendship, and does what is favored by these reasons, but her success in doing so is nonetheless lucky.

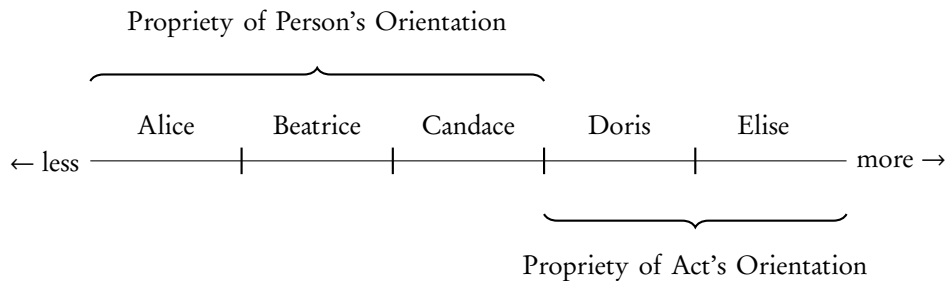
Obviously, *Doris* is not blameworthy here. Nonetheless, her *act* could have been better oriented in one way: its success in conforming to the reasons of friendship could have been a manifestation of her sensitivity to reasons of friendship. While there may be nothing wrong with Doris's orientation to value, we can call her act less than fully proper in its orientation to the reasons of friendship. Compared with the act of Elise—who gets Jim the film because it features Bowie, whom she knows he loves—it is in one way less than ideal. This is not to say that the action is open to strong criticism. But it is increasingly recognized in the

²⁰See especially Harman (2011).

²¹See Parfit (2011), Schroeder (2007, 2008, 2009), Lord (2010), Vogelstein (2012), and Whiting (2014) for discussion of the distinction between apparent and objective normative reasons. Note also that I am using 'sensitive' in its ordinary, dispositional sense, not in Nozick's counterfactual sense.

literature on moral worth that if an agent successfully does what the reasons in fact support accidentally, then even if the agent manifests sensitivity to reasons, there is a respect in which her act could be worthier: its *success* could manifest this sensitivity.²²

These examples illustrate a spectrum of propriety in orientations to value:



It is, I grant, more natural to reserve the expression ‘way of valuing’ for a person’s orientation to value. Nonetheless, I think it is important to recognize the rightmost two points on the spectrum, given that the Extended Hurka Principle applies not only to ways of valuing but also to their manifestations. It is plausible that an act that manifests a proper way of valuing is in that respect non-instrumentally good. It is also plausible that an act that manifests a way of proper valuing and succeeds in doing what the value requires is non-instrumentally better, and better still if its very success manifests that way of valuing. Moreover, while the distinction is admittedly fine, it does seem right to distinguish between how well a *person* is oriented toward a value when she acts, and how well her *acts* are oriented toward value.

So far I’ve been loose in the language I’ve used for the different points on this spectrum. This looseness is unsurprising: ordinary language does not contain terms that exactly circumscribe the boundaries. Nonetheless, for the sake of stating a more precise view, it would be useful to have some terms for the different points on this spectrum. Accordingly, I’m going to regiment the language a little. I’ll call the way of valuing that Alice manifests *commitment*, the way of valuing that Beatrice manifests *weak respect*, and the way of valuing that Candace manifests *strong respect*. I’ll also say that an act which is in conformity with what a value recommends *complies* with value when its conformity manifests strong respect.

The rationale for ‘commitment’ is clear enough: because Alice is following her conscience and trying to conform to the reasons of equality, she shows some commitment to equality. Mere commitment falls far short of what values demand, as we’ve seen. One can be committed to a value while being mistaken, and indeed irrationally so, about what that value requires. The rationale for ‘respect’ is less obvious, because talk of respect can be used to pick out many different things. As Darwall (1977) pointed out, it is important to distinguish between respect as *esteem* and respect as *heedfulness* (which he calls ‘*recognition respect*’). It is the latter sort of respect I have in mind. Darwall originally glossed this kind of respect as follows: ‘There is a kind of respect which can have any of a number of different sorts of things as its object and which consists, most generally, in a disposition to weigh appropriately in one’s deliberations some feature of the thing in question and to act accordingly’ (Darwall 1977, p. 38).²³

Darwall’s gloss on recognition respect is not entirely helpful: what is it to weigh the

²²See, e.g., Arpaly and Schroeder (2014), Mantel (2013) and Lord (2015).

²³In more recent work, Darwall uses ‘recognition respect’ in a narrower, second-personal sense; see esp. Darwall (2006). But the original sense was broader.

feature *appropriately*? Here I think we can distinguish a more subjective notion of appropriateness from a more objective notion, and hence distinguish between *weak* and *strong respect*, as I've suggested. An act manifests weak respect for a value when it manifests a disposition to do what seems likely, relative to the appearances, to correctly respond to that value. It manifests strong respect for a value when it manifests a disposition to what is objectively likely, relative to the evidence, to correctly respond to that value.²⁴

Since I will be defining proper valuing of accuracy in terms of certain positive relations to the norm of accuracy, it must be recognized that the three relations of commitment, weak respect, and strong respect can also hold between a person and a norm. Viewed in this way, they have the following common core: all three ways of properly orienting oneself toward a norm are dispositions to do what the epistemic reasons suggest is likely to be in conformity with the norm. The differences simply turn on how 'epistemic reasons' and 'likely' are understood in this core formula.

To see what I have in mind, consider the familiar distinction between objective and apparent reasons. Objective reasons are facts that count in favor of acts and attitudes by bearing on their correctness. These facts are not all ones to which we have privileged access. Examples in epistemology include the kind of evidence to which we refer when we say things like: 'There was evidence all along that Jones did it. We just discovered it.' Objective reasons are by themselves rarely relevant to justification and rationality. Only apparent reasons and objective reasons that are apparent matter for rationality and justification.

While there is considerable disagreement about the nature of apparent reasons, it is clear that there are such things. Indeed, there are *several* such things, since appearance talk can be used in several ways. Some appearances are belief-relative, others are not. The appearance at issue in the Müller-Lyer illusion is not essentially belief-relative. Call this kind of appearance *seemings-relative*. These appearances are non-factive, but there are also factive appearances, exemplified by *seeing that P*, *remembering that P*, *intuiting that P*, etc. Reasons that are apparent in this sense are objective reasons. Call them *factively apparent reasons*. They can include introspective appearances, like the appearance associated with being aware of some experience. So, even in cases of illusion in otherwise dependable worlds, one has some factively apparent reasons to believe worldly propositions. Facts about one's experiences are generally reliable indicators of external reality in those happy worlds.

With these distinctions in mind, I suggest that for any norm N,

commitment to N in ϕ -ing is grounded in a disposition to ϕ only if, conditional on the belief-relative (*de dicto*) apparent reasons, ϕ -ing is subjectively likely to conform to N;

weak respect for N in ϕ -ing is grounded in a disposition to ϕ only if, conditional on the seemings-relative apparent reasons, ϕ -ing is epistemically likely to conform to N;


strong respect for N in ϕ -ing is grounded in a disposition to ϕ only if, conditional on the factively apparent reasons, ϕ -ing is objectively likely to conform to N;

compliance with N in ϕ -ing is partly grounded in a disposition to ϕ only if, conditional on the all objective epistemic reasons, ϕ -ing objectively likely to conform to N

and fully grounded in conforming to N *by manifesting* that disposition

²⁴The appeal to objective likelihood of compliance doesn't suggest that we are giving an Instrumentalist account of respect. To agree that X-conducivity is necessary for some valuable property is not to hold that this property's value is wholly or even partly explained by X-conducivity. Indeed, one could reverse the order of explanation: X-conducivity matters because it is constitutive of some more fundamental value (e.g., respect).

To put all of this in another diagram:



These things can come apart in more than just the obvious ways. Consider Huck Finn. He believes he is acting wrongly in helping Jim. But his act is praiseworthy all-things-considered. This is because while it fails to manifest commitment to moral norms, it in fact manifests strong respect, and indeed compliance, with these norms. Similar phenomena appear in the epistemic case. Imagine Buck, who has been presented with numerous skeptical arguments that he takes to be decisive, but whose first-order beliefs are unaffected and continue to be responsive to the first-order evidence. Buck is intuitively just as praiseworthy in continuing to hold his first-order beliefs as Huck is in continuing to help Jim. Of course, this is not to say that manifesting a commitment to a norm has no value. As Arpaly (2003) agreed, Huck’s acts would be in one respect admirable if they manifested commitment. All these cases show is that respect is more proper and also more valuable than commitment, not that commitment is not proper or *pro tanto* valuable at all.

3.2.2 Epistemically Proper Valuing in Disposition and Belief: A Special Case

I am now in a position to explain more precisely how I think the Veritist is in a position to make use of the Extended Hurka Principle to explain the value of coherent belief, rational belief, justified belief, and knowledge. Notice first that if we replace ‘N’ with the accuracy norm AN in the earlier general accounts of commitment, respect, and compliance, we get the following characterizations of some ways of valuing accuracy in thought. One manifests:

commitment to AN in believing P iff one manifests a disposition to believe P only if, conditional on the belief-relative (*de dicto*) apparent reasons, P is subjectively likely;

weak respect in believing P iff one manifests a disposition to believe P only if, conditional on the seemings-relative apparent reasons, P is epistemically likely;

strong respect in believing P iff one manifests a disposition to believe P only if, conditional on the factively apparent reasons, P is objectively likely;

compliance with AN in believing P only if (i) one strongly respects AN in believing P, and (ii) the accuracy of one’s belief that P manifests this respect.

These characterizations align with independently attractive views about rationality, justification, and knowledge.²⁵ It is plausible that rational belief is belief in accordance with the apparent epistemic reasons, that justified belief is belief in accordance with objective reasons to which one has access, and that knowledge is accurate belief whose accuracy manifests a disposition to heed objectively good epistemic reasons. Phenomenal conservatives have offered such an account of rational belief.²⁶ Disjunctivists and others have offered such an account of justified belief.²⁷ While there isn’t an existing account of knowledge that aligns perfectly

²⁵I assume that rationality and justification are distinct properties, as Littlejohn (2012) argues. This assumption is not essential for the success of my project, however. I’m just showing that if the two are distinct, we can explain the difference in epistemic value between the two.

²⁶See Huemer (2007).

²⁷See Pritchard (2012) and Littlejohn (2012).

with the view just mentioned, it is attractive in merging themes from virtue epistemology and reasons-first epistemology. Schroeder (2015) argues that knowledge consists in belief for sufficient objective and subjective reasons. If one holds that proper basing on objective reasons simply consists in aptly manifesting a disposition to heed objective reasons, it is natural to see the aforementioned characterization of knowledge as a spinoff of Schroeder's view.

Accordingly, it is plausible, I think, to uphold the following connections:

one's belief that P is *structurally rational* iff it manifests commitment to AN;

one's belief that P is *substantively rational* iff it manifests weak respect for AN;²⁸

one's belief that P is *justified* iff it manifests strong respect for AN;

one's belief that P is *knowledge* iff one complies with AN in believing that P.

There is truth-oriented unity here: coherent, rational, justified, and knowledgeable belief all manifest different positive relations to the norm of accuracy.

Before deriving the epistemic value of these properties using the Extended Hurka Principle, it is worth briefly addressing a concern that I will take up again in §4.2. One might worry that the explanation I'm seeking confronts a dilemma. In particular, one might worry that either the connections I've proposed between the central epistemic properties and the ways of valuing accuracy are implausibly strong, or the way I'm understanding ways of valuing accuracy are implausibly weak. Fortunately, I believe there is a simple response.

In reply to the first horn, I would remind the reader that when the ways of valuing accuracy are understood as I've understood them, these connections are no less plausible than some familiar accounts of rationality, justification, and knowledge. Of course, it is beyond the scope of this paper to defend these analyses of rationality, justification and knowledge. What matters is that even if these analyses are not uncontroversial, it is plausible that there are important epistemic desiderata carved out by the analysantia. There may be other plausible ways of thinking about, for example, knowledge. Indeed, I think there is a kind of knowledge had by animals that doesn't require reasons. But I think it is much less obvious that this kind of knowledge is necessarily better than true belief. Moreover, the more demanding kind of knowledge is not so demanding as to be rare: unsophisticated knowers, including children, often believe things for good epistemic reasons.

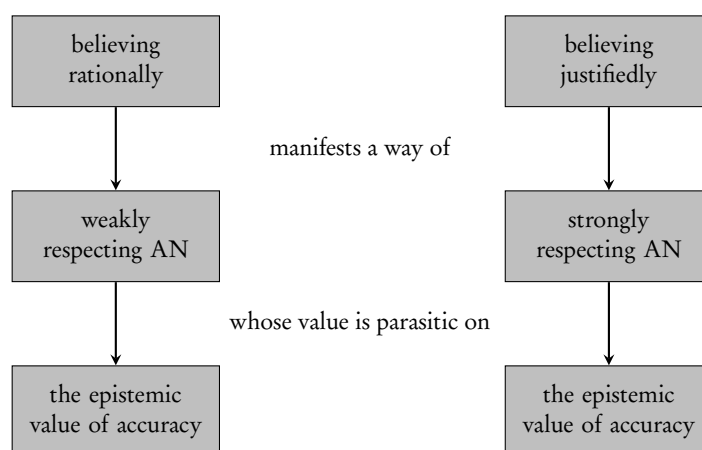
So much for the first horn of the dilemma. And I've already in effect answered the second horn. As I pointed out earlier, the core intuition behind Hurka's Principle extends far and covers valuing in a broad, psychologically undemanding sense. One needn't understand valuing in the thick sense at issue in, say, Watson (1975) for it to be plausible that valuing the good is good. Indeed, one needn't understand valuing as something occurrent, explicitly voiced, emotional, or the like, for it to be plausible that valuing the good is good. Think again of Huck Finn. It is plausible that in acting as he did, Huck manifested proper valuing of Jim's humanity, and that his act was in this respect highly non-instrumentally admirable. This is in spite of the fact that Huck was not only unaware that he was correctly responding to the real reasons, but explicitly looked unfavorably upon what he was doing.

3.3 Deriving Epistemic Value via the Extended Hurka Principle

Once we see that these central epistemic values reflect different ways of valuing accuracy in thought, we can use the Extended Hurka Principle to explain their epistemic value from an

²⁸Ethicists distinguish structural and substantive rationality; cf. Broome (1999) and Kolodny (2005). Structural rationality consists in coherence. Substantive rationality requires heeding good reasons.

accuracy-oriented point of view. Consider rationality. It requires sensitivity (in a dispositional sense) to the apparent evidence bearing on the truth of one’s belief. That is a way to value accuracy in thought. If so, we can use the Extended Hurka Principle to explain the epistemic value of rational belief in accuracy-oriented terms. A similar story would explain the value of justified belief. Putting both in a picture:



A similar story explains why coherence and knowledge have derivative epistemic value. In complying with coherence norms, one manifests commitment to AN.²⁹ In knowing, one complies with AN. These are epistemically proper ways to value accuracy in thought. So we can explain the epistemic value of coherence and knowledge from the perspective of accuracy just like we explained the epistemic value of substantive rationality and justification.

There are epistemic values other than coherence, rationality, justification, and knowledge. But I hypothesize that in a familiar sense of ‘properly epistemic’, every properly epistemic value that is not just instrumentally valuable relative to the others that I discussed will be explicable in something like this framework. In §4, I will address objections to this hypothesis. For now, I am happy to have derived the epistemic value of four central items.

3.4 Comparative Propriety and Rankings of Value

Commitment, weak and strong respect, and compliance are all proper positive relations to AN. If so, coherent belief, rational belief, justified belief, and knowledge are all derivatively good from the perspective of accuracy. Yet one would like to rank these epistemic values. Coherent belief, rational belief, justified belief, and knowledge are not equally epistemically

²⁹One can be accidentally coherent, and so coherent agents can fail to manifest commitment to AN. The claim here is only that *compliance* with requirements of epistemic coherence manifests commitment to AN. Some might doubt the claims about coherence on the grounds that being incoherent can be a way of valuing accuracy: perhaps, for example, one believes claims that one knows cannot all be true, since it is the best one can do by AN’s lights. But first of all, the best examples of what the objector has in mind are the kinds of cases that generate the preface paradox, and in the literature on coherence, it is not uncontroversial that these examples actually involve violations of coherence requirements. Easwaran and Fitelson (2015), for example, argue that deductive consistency is not a requirement of coherence, partly on the basis of considerations about the preface paradox.

A second point is that even if preface cases are understood as the objector suggests, they aren’t counterexamples to the claim on offer. The claim is that in being (non-accidentally) coherent, one manifests proper valuing of accuracy, and that coherence is *pro tanto* derivatively epistemically good for this reason. This claim doesn’t not imply that being incoherent is inconsistent with valuing accuracy. All that follows is that for any agent, if that agent is (non-accidentally) coherent, she manifests one way of valuing accuracy. This second claim doesn’t imply that for any agent, if that agent is incoherent, she does not properly value accuracy in any way.

good. One wants to say that knowledge is better than justified belief, justified belief is better than substantively rational belief (if one draws this distinction), and substantively rational belief is better than structurally rational belief. How should we explain these comparisons?

Although commitment, weak respect, strong respect, and compliance are all proper positive relations to norms, they are not equally proper. For any norm in any domain, it is most proper to comply with the norm. If one must fail to comply, it is more proper to do so with strong respect than without. If one must fail to strongly respect a norm, it is more proper to do so with weak respect than without. If one must fail to weakly respect the norm, it is more proper to do so while at least continuing to manifest commitment to the norm. These general differences explain why

knowledge > justification > substantive rationality > structural rationality

After all, for any norm N, compliance is a more proper relation to N than respect of either kind, and respect of either kind is a more proper relation to N than mere commitment. Since what is best from the perspective of N co-travels with what is most proper to N, we can say that knowledge is best from the perspective of AN. So, compatibly with accuracy's being the *fundamental* epistemic value, we can agree that knowledge is epistemically *best*.

This is not paradoxical. Compliance is the most proper response for any system of norms. This does not mean that there is a further norm of compliance. In addition to a given norm X, one does not need another norm that says, 'Comply with X!' There are just norms, and it is a conceptual truth about norms that compliance is more proper than either kind of respect and either kind of respect is more proper than mere commitment. In the epistemic domain, the fundamental thing is AN. Just as compliance is not a further fundamental norm, knowledge is not a further fundamental value.

Is this illuminating? Yes. We can see the ranking of epistemic value as a special case of a general propriety ordering for positive relations to norms. So, the account is far more illuminating than some view on which it is a bedrock truth that knowledge is better than substantive rationality and substantive rationality is better than coherence.

Have I presupposed what I was trying to explain? No. I appealed to general truths about the comparative propriety of certain positive relations to norms, not particular truths about the comparative goodness of the states that manifest these positive relations to the norm of accuracy. And the appeal to truths about comparative propriety wasn't circular: propriety is a deontic concept, not an evaluative one, and it is non-trivial to draw conclusions about comparative goodness from comparative propriety judgments.

3.5 Advantages

The version of Veritism I've sketched has many attractions. Let's consider them.

My Veritism is the first to honor the fact that value derivation need not proceed *via* instrumental relations. This yields advantages. The Swamping Premise does not generalize to other forms of derivative value. Just as acts that manifest ethically proper ways of valuing ethical values are ethically better, so doxastic attitudes that manifest epistemically proper ways of valuing accuracy are epistemically better.

My view inherits the broader virtues of Veritism. Veritism is a simple axiology. Parsimony is as much a virtue in axiology as it is elsewhere. Supplemented with more sophisticated views about derivative value, Veritism also has great explanatory power. And its explanatory power is intuitive. It is intuitive that we care about things like rationality and justification because we care about accuracy.

Since I deny that the ‘because’ must be understood instrumentally, I can vindicate this intuitive claim more directly than Instrumentalists. Forming beliefs in a way that appears likely to achieve accuracy does not entail reliable achievement of accuracy. Just consider demon worlds, where our experiences are the same but the demon ensures that they radically fail to match external reality. These worlds do suggest that some epistemic values cannot be explained in terms of truth-conducivity. But this does not show that not all epistemic values are truth-oriented. It just shows that truth orientation isn’t a merely instrumental matter.

Isn’t rationality a way of ‘subjectively promoting’ accuracy? Sure. But Instrumentalism makes it unclear why this is more than apparently good. Apparent goodness may be linked to obligation. But it is not a kind of goodness, just as apparent wealth is not a kind of wealth. Yet rationality doesn’t just have apparent epistemic value. It has real epistemic value, even in demon worlds. The only way to explain this short of an error theory is to see ‘subjective promotion’ of accuracy as constituting something further: a way to value accuracy.

4 Six Objections Answered

A theory can have many virtues while facing conclusive objections. Having sketched my theory and explained its virtues, I will now answer six major objections. Of course, I cannot pretend to answer all possible objections to Veritism in this paper. My focus will be on objections that bear on the success of my defense of Veritism against the swamping problem, and of the particular Non-Instrumentalist version of Veritism I’ve sketched.

4.1 What about the Costs of Rejecting Instrumentalism?

One might agree that Instrumentalism has the costs I’ve suggested while worrying that rejecting it will be at least as costly. More specifically, one might think that Instrumentalism does important work for the Veritist that cannot be done in its absence. The Instrumentalist can, after all, tell a helpful story about why we properly care about justification, rationality, and other central epistemic statuses: beliefs with these properties are more likely to be true, and so they help in one’s pursuit of the fundamental epistemic goal. One might think that is not, by contrast, clear that Non-Instrumentalist Veritists can tell a similarly helpful story about why we should care about such statuses. Indeed, why think that caring about accuracy gives one any reason to care about ‘valuing accuracy in thought’ *independently* of whether doing so helps to enhance the accuracy of one’s doxastic inventory?³⁰

Reply. To this objection I’ll offer three overlapping responses. Firstly, the objection overlooks the fact that rejecting Instrumentalism is compatible with accepting all the derivative value and all the axiological explanations that the Instrumentalist accepts. Instrumentalism’s opponent simply insists that there is *more* derivative value around—and more axiological explanations to be given—than Instrumentalists suggest. So rejecting Instrumentalism needn’t carry an explanatory cost. Secondly, in virtue of explaining the value of central epistemic statuses like justification and rationality, my view *does* explain why it is proper to care about them. Moreover, it doesn’t follow from my view—as the objection suggests—that one ought to care about accuracy *independently* of caring about the promotion of accuracy, at least not in any objectionable sense. Finally, Instrumentalism simply doesn’t give a satisfactory answer to the question of why we properly care about all derivative epistemic statuses, if

³⁰Thanks to an anonymous reviewer for pressing this objection.

‘properly’ is understood epistemically.³¹

To elaborate on the first point, let’s consider what rejecting Instrumentalism involves. Instrumentalism claims that all derivative epistemic value is instrumental. Its negation is equivalent to the claim that some derivative epistemic value is not instrumental. This is not the claim that no derivative epistemic value is instrumental. Instrumentalism’s opponents can agree that everything Instrumentalists deem derivatively valuable is valuable *in part* for instrumental reasons. They simply add either that some things are derivatively valuable for multiple reasons (some non-instrumental) or that there are more derivatively epistemically valuable things than Instrumentalists believe. I add both claims. Evidence-sensitive dispositions are epistemically good partly because they are dispositions to hit the mark of truth in actuality, but also partly because they constitute respect for accuracy.³² There are also epistemically good things that Instrumentalism doesn’t dub good. False beliefs that manifest respect for accuracy are epistemically good, as are evidence-sensitive dispositions in the demon world. Given the Dud Principle, Non-Instrumentalist Veritism doesn’t explain the former claim. Since it seems analytic that something cannot have instrumental value in a world *w* if it isn’t instrumental to value in *w*, it also doesn’t explain the latter claim.³³

So as I understand it, Non-Instrumentalist Veritism has all the explanatory power of Instrumentalist Veritism *and more*. Of course, explanatory power isn’t the sole theoretical virtue. And although it is monistic about fundamental value, my view is pluralistic about value derivation relations, and hence less simple than Instrumentalist Veritism. But this loss in simplicity is justified on explanatory grounds: whereas the Instrumentalist Veritist must invoke an error theory,³⁴ the Non-Instrumentalist Veritist can take our intuitions about the value of knowledge, justified belief, rational belief, etc., at face value. It is anyway not clear that reasonable parsimony requires minimizing kinds of derivative value or value derivation. We already acknowledge all the ingredients necessary for constructing these things. It is obvious that there are more relations to fundamental value than promotion (e.g., I value persons, though not primarily by trying to create them), and also clear that there are more axiologically relevant such relations. It is a further question whether there are several *fundamentally* axiologically relevant relations. But I’ve already defended my answer to that question. And again, many people outside epistemology already think there are several types of derivative value. The proposal is simply to take advantage of tools already being used outside epistemology, not to invent new ones.

To be sure, one might wonder whether we must bother with these tools, and whether they really help to explain why we properly care about rationality, justification, and other derivative epistemic values. But the answers to both questions are affirmative. I’ll articulate these answers in turn, showing first why it is wrong to think that Instrumentalism explains why we properly care about the various derivative epistemic values.

Instrumentalism doesn’t plausibly explain our intuitions about the epistemic value of rationality. The inhabitant of the demon world is intuitively *better off* from an epistemic

³¹Perhaps the only thing of non-epistemic value about these statuses is that they are means to accurate belief. But the focus here is on their epistemic value. There is a difference between being an epistemic item that is valuable and being epistemically valuable. Theories of the two may drift far apart.

³²Unless derivative value cannot be overdetermined, we are not barred from making this claim. And surely such overdetermination is possible: the healthiness and pleasantness of some food might each ground sufficient reasons for valuing it, and hence each provide sufficient explanations of why it is valuable (i.e., worth valuing).

³³For a longer defense of this claim, see Sylvan (Ms).

³⁴There are sophisticated error theories in the literature. For example, in addition to the type instrumentalist story critiqued earlier, Goldman and Olson (2009) offer a ‘value autonomization’ explanation of why we *treat* knowledge as more valuable than true belief (consistently with it not really being better). I don’t deny that this error theory has some plausibility, but I prefer to avoid error theories when possible.

point of view *in that world* in virtue of believing rationally. Instrumentalism hence doesn't explain—given that value coincides with that which it is fitting to value—why it is fitting to care about rationality. Perhaps the Instrumentalist will simply deny that it is fitting to value rational belief in the demon world. But this is a significant cost. By contrast, my view offers a simple and appealing story: rationality is valuable even in the demon world because it is constitutive of respect for accuracy, which is a fitting response to the value of accuracy, and hence derivatively valuable via Hurka's Principle. Since it is derivatively valuable, it is fitting to value it. Perhaps the objector doesn't see why we should think that respect is a fitting response to the value of accuracy. But I've already given the argument for that claim: the lesson of the swamping problem and of the demon world *just is* that there are more ways to properly value accuracy than to produce it, and these include respect and commitment.

Now, the question 'Why care about rational belief?' could be heard in a more demanding way, asking why we ought to care about rationality, not just why it is good and hence fittingly valued. But I don't find it obvious that we *ought* to care about rationality, just that doing so is fitting. It is an empirical question whether caring about rationality will itself enable one to respect the norm of accuracy. One might do better to just think hard about interesting questions and respond to one's evidence without intellectual navel-gazing. This is consistent with my view, which only enjoins one to manifest respect for accuracy. Thinking hard about questions and responding to one's evidence can suffice for satisfying that injunction, as I'll stress again in addressing the next objection.

A final challenge still needs to be addressed. It is tempting to agree with the objector that we shouldn't care about accuracy independently of whether doing so helps to enhance the accuracy of one's doxastic inventory. But there are two claims that one could be making here. One could say that it is improper to care about accuracy *without caring* whether one's doxastic inventory is more rather than less accurate. That claim is true, but its truth doesn't support Instrumentalism. Such a combination of care and apathy is *incoherent*, and the opponent of Instrumentalism has an explanation of the badness of this incoherence: one cannot really respect accuracy while not trying to do what by one's own lights would help one to abide by the norm of accuracy.³⁵

There is something else one could be saying in claiming that we shouldn't care about accuracy 'independently' of whether doing so helps to enhance the accuracy of one's doxastic corpus. One could be claiming that it is not fitting to have the preference that one manifest respect for accuracy even when—*unbeknownst to one*—this respect makes no difference to one's accuracy. Given the link between fitting attitudes and value, this claim is equivalent to the claim that it is not epistemically better for the demon-worlder to heed her subjective evidence (which constitutes her manifesting respect for accuracy). Perhaps the Instrumentalist is happy to make this claim. I am not. There is a respect in which the demon-worlder is better off in her own world in virtue of believing epistemically rationally, where 'better off' means *really better off*, not just *better off from her perspective*. Instrumentalism suggests otherwise, and this is a significant cost.

Let's take stock. As we saw, rejecting Instrumentalism has no explanatory cost, and its

³⁵Note that there is no extensional difference between what a (proposition-wise) *subjective* instrumentalist enjoins us to do and what my view enjoins us to do in the demon world: manifesting respect for accuracy is necessarily equivalent to doing what seems most likely to get one the correct answer to the question one is asking. But the fact that the two coincide doesn't suggest that the instrumentalist is right. The important question is why it is genuinely a good thing, and not merely apparently a good thing, to do what is subjectively right. It is easier to answer this question if one takes the value of respect for accuracy to be prior to the value of doing what seems likely to abide by the norm of accuracy. There is nothing strange about taking one of two necessarily equivalent properties to be more fundamental (axiologically or otherwise).

loss in simplicity is justified on the basis of its gains in explanatory power. Instrumentalism struggles to explain why we properly care about epistemic statuses that seem to matter, and has no advantage vis-à-vis others. By contrast, my view does explain why we properly care about the various epistemic statuses that seem to matter. It does so without claiming that we ought to care about accuracy *independently* of whether doing enhances the accuracy of one's doxastic inventory, at least when this claim is understood objectionably.

4.2 Too Demanding?

I said coherence, rationality, justification, and knowledge are epistemically valuable in virtue of manifesting ways of valuing accuracy in thought. But must we value accuracy in thought to believe coherently, rationally, justifiedly, or knowingly? Isn't this too demanding?

Reply. This objection is itself nourished by overly demanding assumptions. In particular, it assumes that valuing accuracy in thought is more demanding than it is.

It is plausible that when we think carefully and heed the apparent evidence bearing on the accuracy of our beliefs, we manifest some respect for accuracy in our thinking. This fits under a more general pattern. When we reason carefully and heed the apparent evidence bearing on whether we would conform to a norm N by ϕ -ing, it is plausible that our ϕ -ing exhibits respect for N. Hence, respecting accuracy needn't be psychologically demanding.

This idea is not strained. If Sal thinks carelessly and forms beliefs iff they are comforting, it is natural to fault him for placing insufficient value on accuracy in his thinking. When we do so, we are calling for something more modest than love or passion for accuracy. Love and passion aren't the only forms of valuing. They didn't figure in my account. We needn't view respect, commitment or compliance as constituted by emotions or states with any fancy feel.

They are just ways to be attuned to factors that bear on the accuracy of one's beliefs. Indeed, ways of valuing accuracy could simply be grounded in competences to correctly respond to reasons.³⁶ This is true, at any rate, if competence is understood as epistemically good *because* its exercises constitute ways to value accuracy. This is not to say that exercising any kind of epistemic competence is sufficient for manifesting a fully proper valuing of accuracy in thought. Norman the reliable clairvoyant perhaps manifests some epistemic competence, but I wouldn't want to say that he values accuracy in his thinking in a fully proper way.³⁷ The competences that constitute measurably proper ways of valuing accuracy are competences to correctly respond to accuracy-relevant reasons, where the reasons are either belief-relative, evidence-relative, or objective.

One might suspect the objection could be revived on the grounds that not all justified beliefs or pieces of knowledge are based on reasons. But here it is worth noting that a belief can manifest a reasons-sensitive competence without being based (at least not consciously) on a reason. Sosa put it well:

[N]o human blessed with reason has merely animal knowledge of the sort attainable by beasts. For even when perceptual belief derives as directly as it ever does from sensory stimuli, it is still relevant that one has not perceived the signs of contrary testimony. . . . [R]eason is always at least a silent partner on the watch for other relevant

³⁶Cf. Greco (1999: 289).

³⁷I am happy to grant that he may exhibit *some* positive relation to accuracy with *some* non-instrumental epistemic value. I am perhaps committed to this claim in virtue of my broad use of 'valuing', but do not see this as a problem. The intuition to honor here is comparative, and my view explains it.

data, a silent partner whose very silence is a contributing cause of the belief outcome. (Sosa 1991, p. 240)

What Sosa is noting is that human knowledge requires reasons-sensitive belief. Yet in believing in a reasons-sensitive way, it is plausible that a believer manifests some kind of respect for accuracy. So even with respect to the beliefs that constitute non-inferential knowledge, we can uphold the picture. While it is implausible to make this claim about *animal* knowledge, it is not a datum that animal knowledge is better than mere true belief.

Of course, it is only in a broad sense of ‘valuing’ that a person who is sensitive to accuracy-relevant reasons automatically counts as manifesting a valuing of accuracy. But as I pointed out earlier, Hurka’s Principle applies no less plausibly to valuing in this psychologically undemanding sense. If anything, it applies *more* plausibly to it. Recall Huck Finn. In a thick sense of ‘valuing’, it may be false that Huck manifests a valuing of moral reasons when he helps Jim and in doing so acts against what he takes the moral reasons to be. Nonetheless, Huck is sensitive to the real moral reasons, and in manifesting this sensitivity, he does recognitionally respect these reasons and his act is morally worthy *for that reason*. Such cases suggest that one is more praiseworthy for manifesting this psychologically less demanding kind of respect than one would be for manifesting the psychologically more demanding kind that implies conscious *de dicto* regard. That fact suggests in turn that this less psychologically demanding kind of respect is *more proper*, and so *non-instrumentally better*.

It is worth stressing one final respect in which my view is less psychologically demanding than it might seem. Notice that my view does not require that for a person to be rational in believing P, this person must have a general disposition to respect the norm of accuracy for any Q she considers and manifest this more general disposition in believing that P. While the characterizations of commitment and respect I offered in §3.2.2 were dispositional, the dispositions were *local* and *proposition-relative*. I said, for example, that one manifests

weak respect for AN in believing P iff one manifests a disposition to believe P only if, conditional on the seemings-relative apparent reasons, P is epistemically likely.

One can manifest this local kind of respect for AN in believing P without manifesting any more global respect for AN. While a *person* will be less admirable if her respect is merely local, her *belief* need not be less admirable so long as it manifests local respect. The point is familiar from some of Hurka’s other work.³⁸ As Hurka notes, virtue evaluations of acts are local and don’t require backing by global traits. If I help a person in need for that person’s own sake, my act is virtuous even if it fails to manifest some more general habit.

Relatedly, it is irrelevant to the negative local evaluation of a person’s acts or attitudes whether that person has some generally admirable disposition that prevents her from performing similarly negatively evaluable acts on a regular basis. As Hurka also notes,³⁹ even if someone isn’t generally disposed to kick dogs for the heck of it—and indeed generally cherishes dogs—we could still condemn that person for kicking one intentionally on a blue moon. Similarly, even if a person cares greatly in general about forming beliefs only if they are accurate, and is disposed in most cases to recognitionally respect the norm of accuracy in belief, it remains possible for that person to fail to recognitionally respect it in some particular cases. In those cases, it is irrelevant to the evaluation of that particular belief that the person in other cases is generally disposed to respect the norm of accuracy.

³⁸See Hurka (2006).

³⁹Hurka (2006, pp. 70-1).

Indeed, my view goes farther, since it claims only that *epistemically (im)proper* ways of valuing accuracy bear on the epistemic evaluation of beliefs. So even if a person sings high praises of accuracy and desires to have accurate beliefs, this fact is irrelevant from an epistemic point of view if such praisings and desires cannot themselves be evaluated as epistemically proper or improper. Yet there is a familiar sense of ‘epistemic evaluation’ having to do with the constitutive standards for doxastic attitudes on which such things aren’t genuine objects of epistemic evaluation. Of course, for Hurka’s Principle to be extended to the epistemic domain, it is necessary that some things worth calling ‘valuings’ can be epistemically evaluated. But the line is not difficult to draw: the valuings that played a role in my account were *belief-forming dispositions*, and such dispositions can be epistemically evaluated in the narrower sense no less than beliefs.

4.3 Doesn’t the View Then Collapse into Familiar Views?

My reply dovetails with another objection. Doesn’t this reply show that my view collapses into familiar views? Aren’t the properties that subvene my ways of valuing accuracy just familiar ones from other theories of rationality and knowledge?

Reply. This objection misses the point of my view. I can agree that my ways of valuing accuracy are grounded in familiar properties from other accounts of the nature of rationality and knowledge. My view is not about the nature of rationality or knowledge, but about why these properties are epistemically good. What matters for this purpose is that these properties manifest something further—viz., ways to value accuracy.

Here we see what is new. My view captures unappreciated truth-oriented unity in epistemic value by its appeal to the Extended Hurka Principle and ways of valuing on accuracy in thought. These ideas are not familiar in epistemology. Most epistemologists assume without argument that all derivative value must be grounded instrumentally.

One might try to restate the objection: ‘But if ways to value accuracy are grounded in familiar properties, then ways to value accuracy will only be epistemically valuable if these familiar properties are epistemically valuable. But it is your job to explain why these familiar properties are epistemically valuable!’ I would reply that they are epistemically valuable because they constitute ways of valuing accuracy, which in turn derive epistemic value *via* the Extended Hurka Principle. This sounds question-begging only if we are hoodwinked by a fallacy. Consider a beautiful painting. It is grounded in a bunch of atoms. These atoms are not beautiful. If they have value, it is because they ground something else. We should not insist that the atoms must have antecedent value and worry that the painting is trash because we cannot find this antecedent value. Their value can be explained in a trickle-down fashion: they ground the painting and it is valuable, so they are valuable.

4.4 Did We Really Avoid Swamping?

I said Veritists could avoid swamping by appealing to two facts: (i) that manifestations of ways of valuing valuable properties are themselves non-instrumentally valuable, and (ii) that this value remains even when the manifestations have the properties that are valued. Thanks to (i) and (ii), it seems true that if an accurate belief manifests (and in this sense has its source in) the believer’s valuing of accuracy, the belief is epistemically better.

But is there really a general truth backing this claim? Suppose a painter values beauty, and for this reason produces a beautiful painting. Does this painting have any extra value

from the aesthetic point of view, simply in virtue of the fact that it was produced by someone who values beauty? This is implausible, one might insist.

Reply. In thinking about this objection, we must be careful to distinguish between *manifestations* of ways of valuing and *mere products* of ways of valuing.⁴⁰ The extended Hurka principle does not say that just any products of fitting ways of valuing have derivative non-instrumental value, only that manifestations of fitting ways of valuing have such value. Of course, it is not implausible that some features of an artwork could manifest an artist's ideals just like her acts and attitudes do. But the examples that make this thought plausible *support* my view. Imagine two novels that appear to have the same stylistic properties. One has them by accident; it is the work of a hack writer, entirely unaware of the significance of these properties, and who could easily have written garbage. The other novel was written by an artist who appreciates the significance of these properties and includes them for this reason. It is not crazy to claim that the second is a better literary work.

More carefully: it is not crazy to make this claim if we assume that the aesthetic value of a work isn't independent of the artistic intention behind it. One might reject that assumption. But if one does reject that assumption, one is adopting the view that the aesthetic properties of a work cannot be *manifestations* of the artist's ideals: they are mere products of the artist. Accordingly, to the extent that we share the objector's intuition about the case above, we can see that as evidence that it is a case of mere production, not of manifestation. If we embrace it as a genuine case of manifestation, we should reconsider the objector's intuition about her (notably underdescribed) case, given that clear-cut cases of manifestation in the aesthetic domain elicit an inconsistent intuition. Hence the objection can be answered with a dilemma: either the case is not a case of manifestation, in which case it doesn't falsify the extended Hurka principle, or it is, and we should reject the objector's intuition on grounds of consistency with other, clearer examples of manifestation.

It is fair for the opponent of Instrumentalism to emphasize the difference between manifestations and mere products. Avoiding Instrumentalism requires an appeal to relations that are not instrumental. The *product of* relation is an instrumental relation. It is not surprising that mere products of valuable sources don't derive any value. And if we consider clear manifestations of the artist's ideals—painterly choices—it is not implausible that they can derive value. Consider a painter who knows what is beautiful and produces something beautiful deliberately, and another who produces something beautiful only accidentally. The first painter's acts are more admirable.

4.5 Does Hurka's Principle Really Generalize?

While Hurka's Principle is intuitive for value *simpliciter*, one might worry that it doesn't generalize to insulated domains like epistemology and grammar. Sometimes I hear the following version of this worry: 'Suppose I love the rules of grammar. My love has no *grammatical* value. Yet doesn't your generalized version of Hurka's Principle predict otherwise?'

Reply. This is not a counterexample to the generalized version of Hurka's Principle. To see this, let me underline a crucial part of the principle:

When V is a non-instrumental value in some domain, proper ways of valuing V in that domain have some derivative non-instrumental value.

⁴⁰This distinction is increasingly popular in the literature on virtue epistemology; see Turri (2011) and Sosa (2015), who argue that the distinction is crucial for a virtue epistemological solution to the Gettier problem.

Is it grammatically proper to love the rules of grammar? Not clearly. But if not, the Extended Hurka Principle does not predict that love for the rules of grammar is grammatically good. For the principle cannot predict that love for X is valuable in any domain unless it is also true that love is a proper response to X in that domain. So the antecedent of the Extended Hurka Principle is not satisfied, and the objectionable conclusion cannot be derived.

One might object that this response is inconsistent with my earlier claim that admiring the beautiful is good from the aesthetic point of view. But there is, I believe, no inconsistency here. It is plausible as a substantive claim about the aesthetic domain that there are domain-relative reasons to have certain pro-attitudes towards the beautiful and the sublime, and certain con-attitudes towards the ugly. For that reason, I think it is plausible that there is a domain-relative kind of value that these attitudes have. Just as the fitting attitudes analysis of value *simpliciter* says that what is valuable *simpliciter* is what it is proper to value *simpliciter*, so a fitting attitudes analysis of domain-relative value would say that what is valuable from the point of view of some domain D is what it is proper from the point of view of D to value. It is left open as a substantive question for particular domains whether there is some rule in the domain enjoining one to value the values of those domains. Here I believe there are substantive, though not structural, disanalogies between domains. Some domains that have rules do not plausibly have *values*;⁴¹ correspondingly, it is less plausible to claim that these domains contain a rule enjoining one to value anything in the domain. Grammar falls in this category. In this respect, there is a substantive disanalogy between grammar and aesthetics. But the Extended Hurka Principle is consistent with these substantive disanalogies: the antecedent of the principle will just be satisfied for some domains, and not others.

If one dislikes substantive disanalogies, one could repeat the response I gave to the worry about the initial strangeness of saying that admiring the beautiful is good from the aesthetic point of view. One could perhaps just accept the claim that loving the rules of grammar is a good thing from a grammatical point of view. For the purposes of my project, it does not matter which response we pursue. I am merely describing my preferred response. While it is more complicated, I believe the complications are appropriate: while I think all normative domains will exhibit structural analogies, it is not implausible to insist that they might be substantively different, by containing different domain-specific norms of propriety.

What about the epistemic domain? Respect, commitment, and compliance in my sense are epistemically proper positive relations to AN. This is a plausible substantive claim about the epistemic domain, one made plausible by the fact that there are recognizable epistemic criticisms that we can issue for people who, for example, disrespect the norm of accuracy in the sense carved out earlier. If so, I can use Hurka's Principle compatibly with denying that loving the rules of grammar is grammatically good.

The fact that our principle appeals to domain-relative propriety does not trivialize it. Instrumentalism is inconsistent with our principle. So, our principle is clearly non-trivial. Might one insist that "proper" means "valuable" and insist that the principle is circular for this reason? No. Propriety is a kind of rightness. Rightness and value represent different species of normativity. There are long-standing debates about how rightness and value are related. So, Hurka's Principle is not circular in linking propriety and goodness.

4.6 What About Other Epistemic Values?

One might accept my account of the epistemic value of rationality, justification, and knowledge but worry that the approach cannot stretch farther. How, one might ask, can it explain

⁴¹Cf. Foot (1985, p. 34) on etiquette.

richer epistemic values like understanding, open-mindedness, and intellectual courage?

Reply. While I think my account can accord these properties all the epistemic value they can reasonably be claimed to have, I should stress first that my burden is lighter than the objection assumes. My goal has not been to offer a full-dress defense of Veritism, but rather to show how Veritists can avoid the swamping problem by dropping Instrumentalism. I've argued that Instrumentalism, not Veritism, is the root of the problem. If this is true, there is not clearly a stronger reason to think the Veritist cannot explain the epistemic value of understanding or open-mindedness than to think other monist axiologists (e.g., knowledge-firsters) or moderate pluralist axiologists cannot. It cannot be taken as a datum that understanding, open-mindedness, and intellectual courage are *fundamentally* epistemically valuable. Since the Veritist can solve the swamping problem, there is no obvious burden she faces that many other monists and moderate pluralists don't face.

Still, one might like a more direct response. I'll offer a two-pronged one. First, I think these rich properties have a special glow partly by having extra value that is not properly epistemic. Beliefs that exhibit these properties may still be epistemically better than mere true beliefs. But—second prong—my framework *can* explain their properly epistemic value.

As a warm-up to the first prong, note that epistemic axiologists often try to explain too much by not distinguishing two things we could mean by 'epistemic value'. Just as Geach (1956) distinguished between attributive and predicative goodness, we should separate

(a) being *epistemically* good

from

(b) being good *simpliciter* and also being epistemic.

There are parallels. Consider talk of a good chess strategy. This may pick out

(c) something that is good *from the point of view of chess strategy*

or

(d) something that is good *simpliciter* and also a chess strategy.

If your chess strategy will humiliate your opponent, it exemplifies (c) but not (d). Similarly, something might be epistemically superb but of little worth absolutely.

Must the chess expert explain the goodness of chess strategies from the Point of View of the Universe? No. It is similarly hard to see why the epistemologist must explain which epistemic items matter from that point of view. Intuitions voiced in the literature on epistemic value often conflate (a) and (b).⁴² But intuitions about (b) are best left for non-epistemic axiologists to explain. A theory of what is valuable from the epistemic point of view need not explain them, just as the chess strategist need not vindicate chess fanaticism.

With this point in mind, I approach the richer properties the objector mentions with some suspicion. Let's take understanding first. While understanding might contribute especially to our flourishing and we do seem to value it as an end *simpliciter*, it is less obvious

⁴²The original *Meno* problem seemed to concern the value *simpliciter* of knowledge, and some contemporary writers, like Kvanvig (2003), take this to be the central problem. Ridge (2013) compellingly argues that there is a widespread and detrimental conflation of attributive and predicative value in the literature.

why it should have premium significance from the epistemic point of view. Understanding does have considerable epistemic value. But Veritism can explain this intuition. Like Pritchard (2010, pp. 75–6), I think one cannot understand a topic without having many rational true beliefs about it, where these beliefs have a special subject matter (i.e., they concern explanatory connections). So understanding will have all the epistemic value that accompanies those many rational true beliefs. Moreover, this value is distinctive in the sense that its subject matter is distinctive. So it is compatible with my view that understanding can have great, even distinctive, epistemic value. I doubt there are any stronger clear-cut intuitions about the properly epistemic value of understanding. To defend this doubt, let's reflect on some relevant facts about attributive value.

A member M of a kind K has more attributive value than another member M^* of K when M is a *better* K , or is *better qua* K , than M^* . And it is trivial that if M is a better K (or is better qua K) than M^* , M^* is a worse K (or is worse qua K) than M . So, we should ask:

(Q) If S 's beliefs B_i constitute less understanding than S^* 's beliefs B^*_j , is it necessarily the case that B_i are worse *qua beliefs* than B^*_j ?

I don't see that our answer to (Q) must be 'yes'. We may hope to understand the world. But it may be unintelligible. A cognizer's world could turn out to be a bag of loose facts, as in a Humean disaster scenario. If Alf's world does turn out to be a bag of loose facts, Alf's beliefs cannot reasonably be regarded as flawed in comparison with those of someone—say, Ralf—in a more orderly world, simply on the grounds that while Ralf grasps connections there to be grasped, Alf grasps no such connections. The problem in Alf's case lies not in his beliefs, but in the world. This contrasts with our attitude towards beliefs that are radically inaccurate, irrational, unjustified, incompetent, and so on. A belief is always in one respect worse as a belief if it is irrational, unjustified, incompetent, or so on.

In saying that the problem in Alf's case is with the world, not his beliefs, I am not saying that whether a belief is flawed *qua* belief cannot turn on how the world turns out. False beliefs are worse *qua* beliefs. But we can do something about beliefs that fail to match the world, just like we can with beliefs that are irrational, unjustified, or so on: we can revise them. Inhabitants of Alf's world cannot do anything about their lack of worldly understanding. So if their beliefs are in all other respects epistemically flawless, it is implausible to claim that they are still worse in the attributive sense than the beliefs of epistemically perfect denizens of Ralf's world.

So I see no straightforward challenge from the value of understanding, provided we focus on the relative attributive value of cognitive states that exhibit and fail to exhibit this property. Again, this is not to deny that there is often more epistemic value in a set of beliefs that exhibits understanding when the believer is confronting a world that can be understood. But Veritism can explain this surplus epistemic value, I suggested. Perhaps there is a further intuition that it is better to know connections than simply to know facts, and that a piece of knowledge about a connection between two facts is better than a piece of knowledge about three loose and separate facts. But this intuition is only vivid when 'better' is left unqualified. I don't see that it is necessary to explain this intuition by appealing to a difference in properly epistemic value. More plausibly, human beings just assign greater final value *simpliciter* to truths about connections than to mere truths.⁴³

⁴³But whatever might be said about the properly epistemic value of *beliefs* about connections, mightn't it be claimed that a *theory* is better *qua theory* if it has systematic unity? And that is a clear intuition about the attributive value of some broadly epistemic item, isn't it?

Let's turn now to open-mindedness, to which similar reasoning applies. Baehr (2012) walks through several characterizations of open-mindedness. The initial accounts he considers have a social dimension, and include evaluative terms with an ethical flavor, like 'fair', 'impartial', and 'honest'. One account views the open-minded thinker as someone who gives 'a fair and impartial hearing to an opposing belief, argument, or body of evidence', while another views the open-minded subject as someone who assesses 'one or more sides of an intellectual dispute in a fair and impartial way' (Baehr 2012, pp. 143–145). While Baehr observes that there are limits to the sociality of open-mindedness,⁴⁴ his final account retains the ethical flavor of earlier accounts. He characterizes all open-minded thinkers as willing to give distinct cognitive standpoints 'a 'serious' (i.e., fair, honest, objective) hearing or assessment' (Baehr 2012, pp. 151–152).

It seems clear that part of what is praiseworthy about open-minded agents is that they manifest an ethical virtue, though it is an ethical virtue *of thought* and an intellectual virtue in the predicative sense. While a Veritist cannot explain this value, she shouldn't be expected to do so. Open-mindedness has some properly epistemic merit. But Veritists face no difficulty explaining this merit. Baehr's own partial account of when it is epistemically appropriate to manifest open-mindedness is close to my account of what it is to manifest respect for accuracy. Baehr tells us that it is appropriate from the epistemic point of view for S to manifest open-mindedness 'only if it is reasonable for S to believe that engaging in this activity in C may be helpful for reaching the truth' (Baehr 2012, p. 160). While I would rephrase this condition to say 'only if S's seeming evidence suggests that engaging in this activity would likely be in conformity with AN', the difference is slight. Both are consistent with Veritism.

I am more skeptical about intellectual courage. Baehr defines intellectual courage as a disposition to persist in aiming at some epistemic good 'despite the fact that doing so involves an apparent threat to one's own well-being' (Baehr 2012, p. 177). I don't see the relevance to properly epistemic praise of the fact that one's persistence in aiming at some epistemic good is done in spite of a threat to one's well-being. Once again, we must distinguish between epistemic virtues in a predicative sense and epistemic virtues in an attributive sense. Epistemic axiology concerns the latter. Intellectual courage isn't a clear example of the latter.

Accordingly, I think my view attributes to understanding, open-mindedness, and intellectual courage all the properly epistemic value that can reasonably be attributed to them. While these are not the only epistemic values other than coherence, rationality, justification, and knowledge, they are commonly mentioned as challenging Veritism.⁴⁵ I see no obvious threat to my strategy for solving the swamping problem if we go beyond a conservative list

I don't find it plausible that systematic unity is a virtue of a theory qua theory when it doesn't reflect corresponding unity in the world. It might be a virtue of the theory qua literary object, but this is irrelevant. While it is fair to worry that the emphasis on cognitive states is too narrow, some narrowness is needed. There are many more precise representational activities with constitutive standards that don't extend to representation *as such*. One map is better than another *qua* map in virtue of being more helpful for understanding how to get from A to B, for any A and B on the map, but it does not follow that representations in general are better *qua* representations if they satisfy this desideratum. In constructing a scientific theory, we are trying to grasp the causal structure of reality, but scientific theorizing is for that reason not *just* representation. Just as satisfying the specific cognitive desires that maps isn't necessarily a value of a representation *qua* representation, so systematic unity isn't necessarily a value of representation *qua* representation.

⁴⁴Summarizing, Baehr writes: '[D]espite the initial plausibility of such a view, open-mindedness cannot be identified with a disposition to set aside one's beliefs about a particular issue in order to consider some opposing viewpoint or piece of evidence (the conflict model). This is because open-mindedness can be exhibited by individuals that are neutral or undecided about the views or arguments they are considering. Second, we have seen that neither is open-mindedness essentially a matter of making a fair or impartial assessment of one or more sides of a dispute...[, since] it can bear on situations in which there is no disagreement...' (Baehr 2012, p.148).

⁴⁵See, e.g., Axtell and Carter (2008), Carter and Gordon (2014), Kvanvig (2003), and Pritchard (2010).

of epistemic values, provided we focus on the attributive value of that larger list's contents.

5 But Isn't There a Less Radical Solution?

I've displayed the advantages of my solution and answered a host of objections. But even if a solution has many advantages and overcomes direct objections, it is only worthy of pursuit if there aren't less radical solutions. Some might insist that we needn't go to such lengths to solve the swamping problem, and that perhaps we can solve it while retaining Instrumentalism. Virtue epistemology (VE) is salient at this point. Some virtue epistemologists might argue that they have already solved the swamping problem in a way that is consistent with Instrumentalist Veritism. If so, we needn't bother with the measures I've recommended.

This response, I will now argue, is mistaken. VE either fails to solve the swamping problem or succeeds only because it collapses into my approach. In defending this claim, I'll show that any version of reliabilist VE that holds onto Instrumentalism cannot solve the swamping problem. Reliabilist VE can only plausibly solve the swamping problem if reliabilist virtues can be seen as constituting something more than mere instruments of truth-promotion: namely, ways of valuing accuracy. While responsibilist VE fares better in solving the swamping problem, it isn't a genuine alternative to my view: as I suggest below and argue at length elsewhere,⁴⁶ my view can be regarded as a form of responsibilism.

5.1 A Dilemma for Reliabilist VE

Reliabilist VE is the version of VE that is most plausibly consistent with Instrumentalist Veritism.⁴⁷ On this view, beliefs gain epistemic value by manifesting an epistemic agent's competence. While competence here is reliabilist, it is also agent-level in a way reliable processes need not be. Owing to this difference, Sosa's model makes a narrower prediction than the process reliabilist's: only beliefs that manifest the agent's epistemic competence gain extra epistemic worth. This can seem to give Sosa a more principled response to the swamping problem. Apt belief is an achievement, while products of reliable processes needn't be.

Compare this with my view. On my view, beliefs gain epistemic worth by manifesting certain ways of valuing accuracy in thought. Structurally, the views are similar: in both cases, the belief derives epistemic worth by manifesting a truth-connected person-level feature. Substantively, they apparently differ: the feature is a reliabilist competence in one case, and a way of valuing accuracy in thought in the other. But one might think Sosa has the advantage, since his feature sounds less lofty. So why prefer my solution?

I have two responses. First, this question rests partly on a false dilemma. I would view some kinds of competent belief as constituting ways to value accuracy in thought. Since I do not understand ways of valuing accuracy in thought in a demanding way, there is no divide between Sosa and me here. The difference rests not in the demands we make, but in our views of derivative epistemic value. On my view, it is only because some competences constitute ways to value accuracy in thought that their manifestations have derivative epistemic

⁴⁶See Sylvan (forthcoming).

⁴⁷Sosa (2007, Ch. 4) is naturally read as trying to reconcile Instrumentalist Veritism with the claim that knowledge is epistemically better than true belief. Sosa doesn't use the terms 'Instrumentalism' or 'Veritism', but he does identify himself as a 'truth-centered' epistemologist, and describes truth-centered epistemology in much the way I've described Instrumentalist Veritism. While he distances himself from the claim that all derivative value is instrumental value, this is because he is using 'instrumental' more narrowly to refer to production value, not product value.

value. Sosa explains why manifestations of competence have epistemic value without regard to whether they constitute ways to value accuracy in thought.

This brings me to the second response. It is crucial for the success of VE that competences be able to ground ways of valuing accuracy in thought. We can see this by seeing why Sosa's picture leaves us with no convincing way to understand the difference in *epistemic worth* between the first-order beliefs of reliable clairvoyants and those of (e.g.) sighted children. While Sosa's view has some advantages over other reliabilist views, the advantages are axiologically insufficient. Perhaps reliable clairvoyants have animal knowledge. The question is whether Sosa can explain the difference in epistemic worth between a reliable clairvoyant's first-order beliefs and those of sighted children. We will see that Sosa's animal/reflective distinction does not help: there is a first-order difference.⁴⁸

Sosa could revise his account of first-order competence in a way that John Greco has recommended. But the revision is attractive because it turns competences into the sorts of things that plausibly ground ways of valuing accuracy. The revision is not motivated unless viewed as conceding that the reliabilist part of the view does not do the axiological work. So, there is a dilemma: this form of VE either (i) is lacking in explanatory power or (ii) collapses into my view.

5.2 First Horn

Competences for Sosa are reliable belief-forming dispositions, relative to favorable conditions in the actual world. This account predicts symmetry between the first-order beliefs of reliable clairvoyants and the first-order beliefs of sighted subjects. Their beliefs can equally manifest Sosa's competence. If the facts of epistemic value are explained by such competences, the first-order doxastic attitudes of reliable clairvoyants must be equal in epistemic value to the parallel attitudes of sighted perceivers.

But even if we agree that the doxastic attitudes of reliable clairvoyants have some epistemic worth, it is implausible that the degree is the same. We might agree that the clairvoyant knows. We are doing epistemic axiology now. Our objection is just as well expressed as one about the difference in epistemic worth between two instances of knowledge.

Some reliabilists try to deem the clairvoyant incompetent by adding a 'no defeaters' clause. But the best versions of this strategy count as defeaters factors that should look irrelevant from a purely reliabilist perspective—e.g., mere beliefs and appearances.⁴⁹ Like

⁴⁸Mark well that I am not denying that the reliable clairvoyant's beliefs have *some* epistemic worth, but just that saying they have *less* worth than the first-order beliefs of (e.g.) sighted children. Given that I am using 'valuing' in a very broad sense, it would be hard for me to deny that reliable clairvoyants might manifest some minimal way of valuing accuracy in thought, even if they fail to manifest many other ways (e.g., respect). I see no need to deny that their beliefs might have some Hurkan epistemic value. Nonetheless, my view easily allows me to discriminate between the *relative* epistemic worth of their beliefs and those of sighted children: sighted children manifest more ways of valuing accuracy in thought, and also manifest more fitting ways of valuing accuracy in thought, than reliable clairvoyants. The objection I'm pressing against Sosa is not that his view wrongly predicts that the reliable clairvoyant's beliefs have some epistemic worth when they have none, but rather that granting Instrumentalism, it fails to explain why their beliefs could have *less* epistemic worth than the beliefs of sighted children.

⁴⁹Some reliabilist accounts of defeat are supposed to address reliable clairvoyance. But they encounter troubles in modified cases. Consider the alternative reliable process account of defeat. On this account, having a defeater for one's belief that *p* consists in having available to one an alternative reliable process that, if used in addition to or instead of the one actually used, would have led to one's not believing *p*. Goldman suggests that this addresses Norman the clairvoyant, claiming that Norman could reliably reason as follows: 'If I had a clairvoyant power, I should surely find some evidence for this. Since I lack any such evidence, I apparently have no such power' (Goldman 1986, p. 112). As Lyons observes, this doesn't go far enough: '[I]t is not difficult to modify the

Greco, I find this unprincipled: ‘Reliabilism insists on a reliabilist account of evidence in favor of belief. . . . [H]ow can the same theory plausibly understand evidence against belief differently? Such a strategy seems at best *ad hoc*’ Greco (2009, pp. 158–9).

Sosa’s tack differs.⁵⁰ He thinks reliable clairvoyants might have animal knowledge and holds that what they lack is reflective knowledge. Reflective knowledge is understood as apt belief that one aptly believes, where aptness is understood in the same externalist way at both orders. What clairvoyants lack is an externalistically parsed grasp of the status of their first-order beliefs. This gives Sosa’s explanation more unity than the defeat strategy.

But the explanation falls short axiologically. Consider a child seeing a red apple on the table in good light and then judging that there is a red apple on the table. This child might lack the conceptual resources to form a belief about the propriety of her first-order belief. Indeed, she might lack the mental state concept BELIEF and the ability to engage in any second-order reflection. Nevertheless, her perceptual beliefs about the apple are epistemically worthier than those of the clairvoyant 400 miles away who equally reliably forms beliefs about the child’s surroundings. This difference is not at the reflective level.

Some might bite the bullet, claiming that the child’s beliefs are on a par with the clairvoyant’s. But the reply is costly, and it underestimates the problem. Many adults lack much of a reflective stance on their perceptual beliefs and so have little reflective perceptual knowledge. The fact that some are in a position to gain reflective knowledge is irrelevant. By analogy: suppose I believe something for bad reasons, though I am in a position to believe for good reasons. That I can do better is irrelevant to the quality of my actual beliefs. The real difference is not second-order. The story should come at the first order.

5.3 Second Horn

John Greco has proposed an enriched account of first-order competence. But his proposal faces his own objections to liberal defeat strategies unless his view is understood as a version of my view. To see this, consider the details. Greco first suggests that we analyze epistemic responsibility in terms of the intellectual dispositions that one manifests when one is motivated to believe the truth:

S’s belief that *p* is epistemically responsible if and only if S’s believing that *p* is properly motivated; if and only if S’s believing that *p* results from intellectual dispositions that S manifests when S is motivated to believe the truth. (Greco 2010, p. 43)

He then analyzes epistemic virtue in terms of reliability *and* epistemic responsibility:

S’s belief that *p* is epistemically virtuous if and only if both (a) S’s belief that *p* is epistemically responsible; and (b) S is objectively reliable in believing that *p*. (Greco 2010, p. 43)

The unjustified but reliable clairvoyant, Greco suggests, fails condition (a).

But Greco’s explanation confronts his own objection to the reliabilist’s appeal to an unreliabilist kind of defeat. Greco insists that reliabilists cannot treat positive and negative justification-relevant properties differently without a unified explanation of why both matter. He faces a similar question. He *says* that epistemic responsibility is a virtue. But we must ask him *why*. What unified story can deem both objective reliability and motivation

counterexample so that it precludes such a reply. We can just stipulate that Norman is so bad at reasoning about such matters that there is no [alternative] process that is both reliable and available to him. This doesn’t seem to change the intuitive verdict’ (Lyons 2009, p. 124).

⁵⁰I rely on Sosa (1991)’s discussion of clairvoyance. Since then he hasn’t explicitly discussed it.

to believe the truth as epistemic virtues? Greco offers no answer. By contrast, my view provides a simple explanation of why being motivated to believe the truth is epistemically virtuous. It is a way to value accuracy in thought, and it matters from the epistemic point of view *via* the Extended Hurka Principle. This confirms my prediction of convergence.

Indeed, there is not just convergence but collapse. For the best explanation of why the dispositions constitutive of being motivated to believe the truth are epistemically virtuous is that they ground ways of valuing accuracy. There is no general reliabilist explanation. If there were, there would have been no need to revise the account! Thus the second horn: revising the account of first-order epistemic virtue is principled only if underpinned by a view of epistemic virtue like mine. So, reliabilist VE either can't explain all facts about epistemic value or collapses into my view.

5.4 About Responsibilism

Of course, reliabilist VE is not the only kind of VE. There is also the responsibilist tradition. While I think others have revealed serious problems for standard forms of responsibilist VE,⁵¹ I've also argued elsewhere that these objections can be overcome if responsibilists drop their emphasis on character traits.⁵² But even if responsibilism is defensible, this fact would not undermine the main point I've been making in this section. For one thing, the traditional version of responsibilist VE is perfectly consistent with my approach, and indeed fares better than reliabilist VE precisely because, on this version, virtues plausibly constitute epistemically proper ways of valuing epistemic value. For another thing, what I've elsewhere recommended as an alternative to the traditional version is an approach whose picture of the virtues directly aligns with the Hurka-inspired approach taken here. Developed in this way, responsibilist VE coincides with the view offered above, and couldn't be seen as an alternative. While the label doesn't matter much to me, I would be happy to see my view as coinciding with the best version of responsibilism, not an alternative to it.

6 Conclusion

Let's take stock. I argued that Veritists can and should reject Instrumentalism. This is the moral of the swamping problem. We can see that it is by seeing that if Instrumentalism were true, the points that led to the swamping problem would create related problems for other economical epistemic axiologies. This shouldn't convince us that a radical pluralism about fundamental epistemic value is true. It should convince us that Instrumentalism is false.

I showed that Veritism is defensible against the swamping problem and related problems if we reject Instrumentalism. As we saw in §2, the derivative value generated by Hurka's Principle is swamping-proof. A beneficent act that manifests the agent's valuing of beneficence is *as such* better than a similar act performed as a PR stunt. This is plausible despite the fact that the goodness of valuing beneficence is parasitic on the goodness of beneficence.

⁵¹The traditional version—'strong conservative' responsibilism, as Baehr (2012) calls it—has encountered many objections in recent years. Critics like Alfano (2012), Baehr (2012), and Dougherty (2011) have argued that knowledge, justified belief, and rational belief do not require the manifestation of stable character traits, as existing versions of strong conservative responsibilism demand. Yet if responsibilism is made less ambitious—taking a 'weak conservative' or 'autonomous' form—it cannot provide a satisfactory account of the epistemic value of rationality, justification, and knowledge. For if virtues are character traits and rationality, justification and knowledge don't require manifestation of such traits, it is unclear why these properties necessarily have any value as such from the epistemic point of view.

⁵²See Sylvan (forthcoming).

This fact would be useless if we couldn't relevantly employ Hurka's Principle in the epistemic domain. But we can. Central epistemic values like coherence, rationality, justification, and knowledge essentially manifest certain ways to value accuracy in thought. None of this is an overintellectualization. To think otherwise is to implicitly overintellectualize the modest business of valuing accuracy in thought. It is natural to fault subjects for failing to place sufficient value on accuracy in thought, and plausible that people who disregard the evidence are open to criticism *because* they fail to place sufficient value on accuracy in thought.

So, there is a way for Veritists to avoid the swamping problem and provide explanations of the epistemic value of central epistemic properties like rationality and knowledge. Since it is antecedently attractive to think that epistemic evaluation is truth-oriented, I think we should continue to be attracted to Veritism.

Of course, as I've stressed throughout, my aim has not been to give a conclusive defense of Veritism, but rather a defense from the swamping problem and attendant problems. An exhaustive argument for Veritism would require showing that it is preferable to all other economical epistemic axiologies, and that it can explain the epistemic value of anything that might be thrown at it. But while it was beyond the scope of this paper to offer such a comprehensive defense, I do think the central ideas in this paper further aid this more ambitious project. I will conclude by briefly explaining why, beginning with a comparison with other economical epistemic axiologies.

As we saw earlier, no modest pluralism about fundamental epistemic value that clings to Instrumentalism can avoid every relative of the swamping problem. With Instrumentalism dropped, there is no swamping-related reason to prefer these forms of pluralism. Since parsimony is a virtue, we have reason to prefer my view to these views. As we've also seen, other forms of monism that preserve Instrumentalism will face relatives of the swamping problem. If the desire to avoid swamping attracts one to a knowledge-first epistemic axiology, one is short-sighted: this view faces related problems, as I argued earlier. Of course, one could embrace a non-Veritist monism while also rejecting Instrumentalism. Because I am optimistic about the analyzability of knowledge, I doubt that the main alternative—i.e., the knowledge-first view—is well-motivated. With such motivations set aside, I see no clear reason to think that non-traditional monisms will outperform my more traditional monism. But this is not the place to argue against the knowledge-first program. Here I claim only that it offers no clear *axiological* advantage that isn't matched by an advantage of my view.

More radical pluralists might claim that there is such a rich variety of epistemic values that it is *a priori* implausible to expect any economical axiology to succeed. Or they might insist that it is clear that we do value, say, knowledge, understanding, rationality, etc., for their own sake. But my Non-Instrumentalist Veritism offers the intuitive satisfaction that more radical pluralist views offer without losses in unity and simplicity. For it is consistent with my view that there are indefinitely many *non-instrumental* or *final* epistemic values.⁵³ Accordingly, it is consistent with my view that in one sense of 'for their own sake', we properly value many epistemic properties 'for their own sake' (i.e., not as *mere means*). But as I've shown, it is consistent with embracing a richly pluralist inventory of *final* epistemic values to think that there is just one *fundamental* epistemic value. For if derivativeness doesn't imply instrumentality, then non-instrumentality also doesn't imply fundamentality.

Finally, it is worth re-emphasizing that it is consistent with my view that any number of epistemic properties might be finally valuable *simpliciter*. As I've emphasized, we should distinguish being valuable and epistemic from being epistemically valuable, and as epistemic

⁵³In accordance with the etymology ('finis' in Latin means *end*), I assume that 'final' means *as an end, not merely a means*, and hence that final value and non-instrumental value are the same thing.

axiologists confine our attention to the latter. And I believe the response I gave to some other example-based challenges to Veritism in §4.6 will extend to more exotic epistemic properties that might be mentioned, such as intellectual honesty, intellectual perseverance, epistemic justice, contemplativeness, and so on. Admittedly, given that the early literature on epistemic value was concerned with the *Meno* problem and it is far from clear that this is simply a problem about attributive epistemic value, this focus might be narrower than that of some earlier literature. But I don't see that epistemologists will be particularly well-equipped to adjudicate monism/pluralism debates in this wider arena. So I think my line is the most ambitious that could reasonably be expected in an epistemology paper such as this.

References

- Adams, Robert 2006, *A Theory of Virtue* (Oxford: Oxford University Press)
- Alfano, Mark 2012, 'Expanding the Situationist Challenge to Responsibilist Virtue Epistemology', in *Philosophical Quarterly* 62: 223–249
- Anderson, Elizabeth 1995, *Value in Ethics and Economics* (Cambridge MA: Harvard University Press)
- Arpaly, Nomy 2003, *Unprincipled Virtue* (Oxford: Oxford University Press)
- Arpaly, Nomy and Schroeder, Timothy 2014, *In Praise of Desire* (Oxford: Oxford University Press)
- Axtell, Guy and Carter, J. Adam 2008, 'Just the Right Thickness: A Defense of Second-Wave Virtue Epistemology', in *Philosophical Papers* 37(3): 413–434
- Baehr, Jason 2012, *The Inquiring Mind* (Oxford: Oxford University Press)
- Berker, Selim 2013, 'Epistemic Teleology and the Separateness of Propositions', in *Philosophical Review* 122.3: 337–393
- Broome, John 1999, 'Normative Requirements', in *Ratio* 12: 398–419
- Carter, J. Adam and Gordon, Emma C. 2014, 'Open-Mindedness and Truth', in *Canadian Journal of Philosophy* 44(2): 207–224
- Carter, J. Adam and Jarvis, Benjamin 2012, 'Against Swamping', in *Analysis* 72(4): 690–699
- Darwall, Stephen 1977, 'Two Kinds of Respect', in *Ethics* 88: 36–49
- Darwall, Stephen 2006, *The Second-Person Standpoint* (Cambridge MA: Harvard University Press)
- Dougherty, Trent 2011, 'Reducing Responsibility', in *European Journal of Philosophy* 18: 1–14
- Easwaran, Kenny and Fitelson, Branden 2015, 'Accuracy, Coherence, and Evidence' in *Oxford Studies in Epistemology* 5 (Oxford: Oxford University Press)
- Feldman, Fred 2000, 'Basic Intrinsic Value', in *Philosophical Studies* 99: 319–46
- Geach, Peter 1956, 'Good and Evil' in *Analysis* 17(2): 33–42
- Goldman, Alvin 1986, *Epistemology and Cognition* (Cambridge MA: Harvard University Press)
- Goldman, Alvin and Olsson, Erik 2009, 'Reliabilism and the Value of Knowledge' in Haddock, A., Millar, A., and Pritchard, D. (eds.) *Epistemic Value* (Oxford: Oxford University Press)
- Goldman, Alvin 2012, *Reliabilism and Contemporary Epistemology* (Oxford: Oxford University Press)
- Greco, John 1999, 'Agent Reliabilism', in *Philosophical Perspectives* 13: 276–96
- Greco, John 2004, 'How to Preserve Your Virtue While Losing Your Perspective' in Greco, J. (ed.) *Ernest Sosa and His Critics* (Oxford: Blackwell)

- Greco, John 2010, *Achieving Knowledge* (Cambridge: Cambridge University Press)
- Harman, Elizabeth 2011, 'Does Moral Ignorance Exculpate?' in *Ratio* 24(4): 443–468
- Harman, Gilbert 1993, 'Desired Desires' in Frey, R. and Morris, C. (eds.) *Value, Welfare, and Morality* (Cambridge: Cambridge University Press)
- Herman, Barbara 1993, *The Practice of Moral Judgment* (Cambridge MA: Harvard University Press)
- Huemer, Michael 2007, 'Compassionate Phenomenal Conservatism' in *Philosophy and Phenomenological Research* 74(1): 30–55
- Hurka, Thomas 2001, *Virtue, Vice and Value* (Oxford: Oxford University Press)
- Hurka, Thomas 2006, 'Virtuous Act, Virtuous Dispositions' in *Analysis* 66: 69–76
- Jones, Ward 1997, 'Why Do We Value Knowledge?' in *American Philosophical Quarterly* 34: 423–439
- Kagan, Shelly 1998, 'Rethinking Intrinsic Value' in *Journal of Ethics* 2: 277–297
- Kolodny, Niko 2005, 'Why Be Rational?' in *Mind* 114: 509–563.
- Kolodny, Niko 2007, 'How Does Coherence Matter?' in *Proceedings of the Aristotelian Society* 107: 229–263.
- Korsgaard, Christine 1983, 'Two Distinctions in Goodness' in *Philosophical Review* 92: 169–195
- Kvanvig, John 2003, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press)
- Littlejohn, Clayton 2012, *Justification and the Truth Connection* (Cambridge: Cambridge University Press)
- Lord, Errol 2010, 'Having Reasons and the Factoring Account' *Philosophical Studies* 149: 283–296
- Lord, Errol 2015, 'Acting for the Right Reasons, Abilities, and Obligation' in *Oxford Studies in Metaethics* 10 (Oxford: Oxford University Press).
- Mantel, Susanne 2013, 'Acting for Reasons, Apt Action, and Knowledge' *Synthese* 190: 3865–3888
- Moody-Adams, Michelle 1994, 'Culture, Responsibility, and Affected Ignorance' *Ethics* 104(2): 291–309
- Moore, G. E. 1903, *Principia Ethica* (Cambridge: Cambridge University Press)
- Nozick, Robert 1984, *Philosophical Explanations* (Cambridge MA: Harvard University Press)
- Parfit, Derek 2011, *On What Matters* (Oxford: Oxford University Press)
- Pritchard, Duncan 2011, 'What Is the Swamping Problem?' in Reisner, A. and Steglich-Petersen, A. (eds.) *Reasons for Belief* (Cambridge: Cambridge University Press)
- Pritchard, Duncan 2012, *Epistemological Disjunctivism* (Oxford: Oxford University Press)
- Rabinowicz, Wlodek and Ronnow-Rasmussen, Toni 2000, 'A Distinction in Value: Intrinsic and For Its Own Sake' in *Proceedings of the Aristotelian Society* 100: 33–51
- Rashdall, Hastings 1907, *The Theory of Good and Evil* (Oxford: Clarendon Press)
- Ridge, Michael 2013, 'Getting Lost on the Road to Larissa' in *Nous* 47(1): 181–201
- Scanlon, T. M. 1998, *What We Owe to Each Other* (Cambridge MA: Harvard University Press)
- Schroeder, Mark 2007, *Slaves of the Passions*. (Oxford: Oxford University Press)
- Schroeder, Mark 2008, 'Having Reasons' in *Philosophical Studies* 139: 57–71
- Schroeder, Mark 2009, 'Means-End Coherence, Stringency, and Subjective Reasons' in *Philosophical Studies* 143: 223–248
- Smith, Michael 1994, *The Moral Problem* (Oxford: Blackwell)
- Sosa, Ernest 1991, *Knowledge in Perspective* (Cambridge: Cambridge University Press)
- Sosa, Ernest 2007, *A Virtue Epistemology (v.1)* (Oxford: Oxford University Press)

- Sosa, Ernest 2015, *Judgment and Agency* (Oxford: Oxford University Press)
- Sylvan, Kurt ms, 'Respect for the Truth and the Normativity of Epistemic Rationality' (Unpublished Manuscript, University of Southampton)
- Sylvan, Kurt Forthcoming, 'Responsibilism out of Character' in Alfano, M. and Fairweather, A. (eds.) *Epistemic Situationism*. (Oxford: Oxford University Press)
- Turri, John 2011, 'Manifest Failure: The Gettier Problem Solved' in *Philosophers' Imprint* 11(8): 1–11
- Vogelstein, Eric 2012, 'Subjective Reasons' in *Ethical Theory and Moral Practice* 15: 239–257
- Watson, Gary 1975, 'Free Agency' *Journal of Philosophy* 72: 205–20
- Wedgwood, Ralph 2002, 'The Aim of Belief' in *Philosophical Perspectives* 16: 267–297.
- Wedgwood, Ralph 2013, 'Doxastic Correctness' in *Proceedings of the Aristotelian Society Supplementary Volume* 87: 217–234.
- Whiting, Daniel 2014, 'Keep Things in Perspective: Reasons, Rationality, and the A Priori' in *Journal of Ethics and Social Philosophy* 8(1): 1–22
- Williams, Bernard, 1981, *Moral Luck* (Cambridge: Cambridge University Press)
- Zagzebski, Linda 1996, *Virtues of the Mind* (Cambridge: Cambridge University Press)
- Zagzebski, Linda 1999, 'From Reliabilism to Virtue Epistemology' in *Proceedings of the Twentieth World Congress of Philosophy* (Bowling Green, OH: Philosophy Documentation Center)
- Zagzebski, Linda 2003, 'The Search for the Source of the Epistemic Good' in *Metaphilosophy* 34: 12–28.
- Zagzebski, Linda 2004, 'Epistemic Value Monism' in Greco, J. (ed.) *Ernest Sosa and His Critics* (Oxford: Blackwell)
- Zimmerman, Michael 2001, *The Nature of Intrinsic Value* (Lanham: Rowman & Littlefield)