

## **The Value Turn in Epistemology**

Wayne Riggs

### **Introduction**

From the somewhat grandiose title I have chosen for this chapter, the reader may be excused for expecting a proclamation that a revolution in epistemology has taken place, or perhaps a call to arms for the people to rise up and begin one. Certainly, when Richard Rorty wrote the introduction to the famous anthology whose title mine alludes to<sup>1</sup>, he described his subject matter as constituting a revolution, at least in the minds of its proponents. Alas, my aims here are less bold than a call to revolution, but nevertheless I believe there is a small but growing tendency to approach questions in epistemology in a refreshingly different way—a way that brings new resources to bear on old questions, much as the “linguistic turn” brought the resources of logic, linguistics and philosophy of language to bear on the old questions of metaphysics, epistemology, and ethics. Unlike that famous “turn,” however, this way of approaching epistemology (hereafter called “value-driven epistemology”) does not call for the elimination of all other methods. It is largely a complementary, rather than a rival, way to think about epistemological problems. Furthermore, it brings with it a way of conceiving epistemology that is less confining and intellectually richer than the more traditional way that tends to define epistemology as merely the “theory of knowledge.” While value-driven epistemology need not deny the importance or even centrality of knowledge and the pursuit of an analysis of knowledge to epistemology, at the very least it argues for a widening of epistemological interest to other cognitive phenomena.

---

<sup>1</sup> Rorty, *The Linguistic Turn*

In what follows, I will characterize value-driven epistemology in a way that other epistemologists who I would categorize as sympathetic to the cause or even card-carrying members of the value-driven epistemology club might object to. That is due partly to the irascibility of philosophers, but also to the fact that we are only just beginning to think of what we do as having a distinct method, or as embodying a distinct approach, from that of other epistemologists. So consider this an early foray in to the wilderness; a first attempt to make concrete what holds together a certain collection of epistemological views and projects. I will begin by placing value-driven epistemology in the context of methodologies within epistemology more generally. I will then argue for the peculiar benefits it promises to offer epistemology.

### **The Problem of the Criterion (& the Language, & the Values, &tc)**

Though I will be arguing otherwise later, for now let's assume that the main task of epistemology is to articulate the nature and scope of knowledge. In other words, the task is to tell us what knowledge is, and how to tell, given any putative item of knowledge, whether it really is one or not. Roderick Chisholm famously characterized epistemology this way, and then described epistemologists' methodological situation as being beset with what he called the "problem of the criterion."

We may distinguish two very general questions... The first of these may ...be put by asking, "What is the *extent* of our knowledge?" and the second by asking, "What are the *criteria* of knowing?"

If we know the answer to either one of these questions, then, perhaps, we may devise a procedure that will enable us to answer the other. If we can specify the criteria of knowledge, we may have a way of deciding how far our knowledge extends. Or if we know how far our knowledge does extend and are able to say what the things are that we know, then we may be able to formulate criteria enabling us to mark off the things we do know from the things that we do not.

But if we do not have the answer to the first question, then, it would seem, we have no way of answering the second. And if we do not have the answer to the second, then, it would seem, we have no way of answering the first. (Chisholm [1989], 6, emphasis in original)

According to Chisholm, we can make progress on either half of the epistemological task only by assuming that we have already largely accomplished the other half. For example, if we assume we already have a good idea of the criteria of knowledge, then we can use those criteria to judge putative instances of knowledge to see if they meet the criteria or not. Alternatively, if we take our individual judgments about particular cases of knowledge to be generally reliable, then we can generalize from those to formulate some criteria of knowledge. But absent either a set of criteria we can trust or a set of particular judgments we can trust, we can not even begin to determine either the nature or the scope of our knowledge.

Chisholm called the first of these methodological stances “generalism,” and the second “particularism.” Chisholm clearly favored the latter method. A “purist” of either methodological stripe would presumably proceed rigorously in only one direction. That is, a purist particularist

would take all the non-controversial particular knowledge judgments as rock-solid data, and adjust the criteria of knowledge as necessary to account for that data. Similarly, a purist generalist would take the criteria to be a given, and bite whatever bullets were necessary regarding individual knowledge claims. But, of course, there are no extant purists of either stripe. Particularism is the declared method of choice for most epistemologists, but they are clearly no purists. That is partly because the “data” of individual knowledge claims, judgments, intuitions, or what have you, are not obviously consistent. So, some of the data must be given up for the sake of a consistent theoretical account. But what guides this culling of intuitions about particular cases? I think the obvious answer is the epistemologists’ convictions about what the criteria are like. In other words, which data are taken to be signal and which noise will be determined in large part by what sort of knowledge criteria the philosopher who is building the theory finds more palatable.

None of this is meant by way of criticism. I think that a kind of reflective equilibrium that takes seriously the most initially plausible particular judgments as well as initially plausible criteria is the right methodology to follow. No purist, I. But that is just my point. Chisholm was wrong to say that one must choose one or the other methodology. He was right that one cannot provide a completely independent argument for both the particular judgments and the criteria, but that does not preclude us from taking our initial strongest convictions about both and trying to find an account that ties them nicely together.

Nonetheless, it was perceptive of Chisholm to see the way our judgments about particular cases and our conclusions about the criteria of knowledge are systematically interrelated in this quite intimate way. Nevertheless, his “problem of the criterion” represents only a part of the whole picture, because there are other sources of intuitions, principles, and arguments that bear

on the questions of the nature and extent of our knowledge. In this section, I will describe and discuss briefly two of these sources, though I readily admit there are more. One of these sources is our linguistic practices involving the use of the word “know” and its close cognates. The other is our theory of epistemic value.

*The New Ordinary Language Epistemology*

Though the theory called “contextualism” in epistemology has been around for a while, discussions of it have multiplied at an astonishing rate in the last few years. As a consequence, perhaps, what the term means has shifted considerably from its original use. Keith DeRose characterizes its recent incarnations in his “Contextualism: An Explanation and Defense.” He says that

“contextualism” refers to the position that the truth-conditions of knowledge ascribing and knowledge denying sentences (sentences of the form “S knows that p” and “S doesn’t know that p” and related variants of such sentences) vary in certain ways according to the context in which they are uttered. What so varies is the epistemic standards that S must meet (or, in the case of a denial of knowledge, fail to meet) in order for such a statement to be true. (DeRose 1999, p. 187)

On the face of it, as DeRose himself admits a few paragraphs later, this is not a theory of knowledge in any ordinary sense. It is, rather, a theory of the semantics of knowledge ascriptions and denials. Yet DeRose defends the idea that such a semantic theory for “knows” holds a great deal of significance for anyone engaged in the task of determining the nature and scope of knowledge. What is interesting about this project for present purposes is that DeRose explicitly endorses a methodology that begins by paying attention to linguistic data. The best evidence for

contextualism, says DeRose, is the actual (or realistically imagined) linguistic practices of ordinary people when they use the term “knows.”

Contextualists typically appeal to pairs of cases that forcefully display the variability in the epistemic standards that govern ordinary usage: A “low standards” case in which a speaker seems quite appropriately and truthfully to ascribe knowledge to a subject will be paired with a “high standards” case in which another speaker in a quite different and more demanding context seems with equal propriety and truth to say that the same subject (or a similarly positioned subject) does not know. The contextualist argument based on such cases ultimately rests on the premises that the positive attribution of knowledge in the low standards case is true, and that the denial of knowledge in the high standards case is true. Why think that both of these claims are true? (DeRose 2005, p. ???)

After posing the question, DeRose answers it revealingly,

[S]ince the contextualist’s cases do not involve speakers who are involved in some mistaken belief about a relevant underlying matter of fact, there is good reason to think that their claims, which are made with perfect propriety, are true, and it’s a bad strike against a semantic theory if it rules these claims to be false... (DeRose 2005, p. ???)

DeRose accepts a general principle to the effect that the utterances of speakers who are not mistaken about a relevant underlying matter of fact and who speak with propriety in a given

situation, are generally true. This gives him a set of data (the set of presumed true utterances of sentences of the forms ‘S knows that p’ or ‘S doesn’t know that p’) that is independent of any prior assumptions about either the nature or extent of knowledge. Furthermore, DeRose insists that these data, and the semantic theory of “knows” that he develops from them, are importantly relevant to the theory of knowledge. Thus, if DeRose is right, we have a source of intuitions, principles and arguments relevant to the theory of knowledge that do not require that we already have in hand either a criterion of knowledge or a set of judgments about particular cases of knowledge.

The promise of this approach is evident in both the interest it has aroused as well as the variety of positions this interest has produced. Of those who deny contextualism, called invariantists, there are now several varieties, which differ along several different dimensions. This is not the place to summarize that interesting literature, but suffice it to say that explaining the semantics of “knows” turns out to be a philosophically interesting and challenging task in its own right. The biggest limitation I find in this literature is a clear and persuasive explanation of precisely how this semantic dispute is supposed to affect the development of theories of knowledge. There are multiple possibilities, none of which, it seems to me, has been clearly worked out yet.

### *Value-Driven Epistemology*

The other source of intuitions, principles and arguments I want to talk about that bear on the nature and scope of knowledge that Chisholm’s “problem” does not acknowledge is a view about what our epistemic values are. Until recently, there was little critical discussion in contemporary epistemology of what goals, values or ends were constitutive of the normative discipline of epistemology. It was largely taken for granted that the fundamental, and often it was

said, the only, epistemic goal, aim or value was the accumulation of true beliefs. When pressed, epistemologists would often add that we would like to avoid false beliefs as well, almost as though that was just another way of putting the first point. But recent debate about precisely how to specify our epistemic goals has shown that the issues are both trickier, and more fundamental, than had been appreciated before (get refs).

This debate is really just getting started, so we cannot with any confidence point to any real consensus about what the correct set of epistemic goals or values will consist of. But that is not required to make the present point. As we make progress in determining these goals, these results will have consequences for the theory of knowledge. For instance, if we assume a certain set of epistemic values, then we can use the philosophical tools of axiology to narrow down our conception of knowledge. For example, if we decide that having true beliefs is the only basic epistemic value, then any analysis of knowledge must be such that one can trace the value of knowledge through value-transmitting relations to the value of having true belief. Of course, most accounts of knowledge claim that knowledge is a species of true belief, so every analysis would be such that knowledge derives value from the value of having true belief trivially. But, as Jon Kvanvig has argued, that is not the value claim that is usually made in favor of knowledge. Knowledge, qua knowledge, must be more valuable than any subset of its components, or else there would be no reason to value knowledge itself, but just those valuable subparts. Therefore, any analysis of knowledge must be able to show how knowledge, qua knowledge, derives its value from the value of having true beliefs. The fact that one can place substantive constraints on a theory of knowledge based solely on a commitment to the fundamental epistemic values from which any normative epistemic property or phenomenon must derive its value indicates again how Chisholm was wrong. We can get started at least on the project of determining the nature

and scope of knowledge without assuming much in the way of particular judgments in putative instances of knowledge or about the criteria of knowledge. All we need is the claim that  $X_1 \dots X_N$  are the fundamental epistemic values from which everything of epistemic value must be derived.

### **What Is Value-Driven Epistemology?**

So far, all I have argued is that there are multiple “starting points” for the theory of knowledge. In addition to the usual methodology of “particularism,” there is also the ordinary language method of contextualism and its cousins, as well as the method of value-driven epistemology. This alone would give ample reason for there to be epistemologists who pursue this method, since an intellectual division of labor is more efficient than everyone’s following the same path. And that might be all there was to say in favor of value-driven epistemology, if epistemology was no more than the theory of knowledge. But there are some good reasons to suppose that it isn’t. William Alston defends one such broader view of epistemology in his recent book, *Beyond Justification*. One does not have to accept his central thesis that we should abandon theorizing about a single, univocal property of “epistemic justification” to agree with his characterization of epistemology.

What we call ‘epistemology’ consists of some selection from the problems, issues, and subject matters dealt with by philosophers that have to do with what we might call the cognitive side of human life: the operation and condition of our cognitive faculties—perception, reasoning, belief formation; the products thereof—beliefs, arguments, theories, explanations, knowledge; and the evaluation of all that. And so a

very broad conception of epistemology would be *philosophical reflection on the cognitive aspects of human life*... (Alston, 2005, p. 2-3)

I find this characterization of epistemology fairly agreeable, though still a little rough. I will offer a slightly more formal characterization of my own later. But for now, I simply want to point out that one need not think of epistemology as being exhausted by the theory of knowledge. For one thing, it's not hard to imagine someone who has a head full of knowledge, but who would not be what anyone would consider an "epistemic exemplar." Here we can imagine phone-book memorizers and the like. Such (typically imaginary) people can lay claim to huge amounts of knowledge, yet their cognitive life seems to lack much epistemic value. This, I take it, is one indication that offering an account of knowledge is not all there is to do in epistemology.

Furthermore, it should be evident that a great many things can be evaluated from what Alston calls "the epistemic point of view" (Alston, get ref), on which more below. These include beliefs, sets of beliefs, believings, habits of belief-formation, intellectual traits, persons, and so on. It is unlikely that all of these things will be included in any one theory of knowledge. And yet it is indisputable that each of these kinds of thing can have properties that are better or worse from an epistemic point of view. To take just one example, few currently popular theories of knowledge require that one's beliefs be coherent. Yet surely we all would take the coherence of our beliefs to be epistemically desirable. Would not, then, the specification of the conditions of belief coherence be a part of epistemology, even if the coherence theory of knowledge is false?

Once one accepts that there is more to epistemology than the theory of knowledge, one is led immediately to wonder how one determines the limits of the realm. One way to approach

epistemology is to think of it as a set of problems, or puzzles, to solve. There is the problem of knowledge, which for many amounts to the problem of epistemic justification. Under that heading, there are many more problems; the Gettier problem, the lottery and preface paradoxes, the internalism/externalism debate, the skeptical problem, the new evil demon problem, and so forth. There are a few other problems that typically make the list: the problem of rationality, the problem of the a priori, and sometimes the problems of induction. No doubt I have left some perennial favorites off my lists, but my point is that this approach to epistemology does not think of it so much as a domain of inquiry as a set of problems to be solved.

Value-driven epistemology invites, though it might not insist, that one think of epistemology as a normative domain of inquiry—one that is bounded largely by the values that are fundamental to it. Determining these values is itself one of the tasks proper to value-driven epistemology. Indeed, it is one of the most important tasks, since theorizing in value-driven epistemology often begins by appeal to those very values. Furthermore, it is only by appeal to these values that epistemology has a reasonable claim to being a normative discipline. Consider the way William Alston conceives of the normativity of epistemology in his classic paper, “Concepts of Epistemic Justification.”

[O]ne can evaluate S’s believing that *p* as a good, favorable, desirable, or appropriate thing without thinking of it as fulfilling or not violating an obligation...[B]elieving can be evaluated from various points of view, including the epistemic, which, as we have noted, is defined by the aim at maximizing truth and minimizing falsity. It may be a good thing that S believes that *p* for his peace of mind...[b]ut none of this would render it a good thing for S to believe that *p* from the epistemic point of view. To believe that *p*

because it gives peace of mind...may not be conducive to the attainment of truth and the avoidance of error. [Alston ?????]

Here Alston refers to an “epistemic point of view,” from within which we can evaluate whether something is a good thing or not from that point of view. What does this involve? It seems simple. Something is good from the epistemic point of view if it is appropriately related to the values that determine the normative boundaries of the epistemic. In this paper, Alston explicitly identifies these values as being “the attainment of truth and the avoidance of error.” Furthermore, he implicitly identifies the relation that must hold between the object of evaluation and the values in play. He claims that believing that *p* is not a good thing from the epistemic point of view unless it is *conducive* to epistemic values. So the relation that matters here is a means-end relation. To put it all together, Alston is claiming that *X* is a good thing from the epistemic point of view only if *X* is conducive to the attainment of truth and the avoidance of falsehood.

In his more recent book, Alston is more pluralistic, and acknowledges that there might be epistemic values beyond the attainment of truth and the avoidance of error. But whatever one takes the fundamental epistemic values to be, this idea of the “epistemic point of view” gives us a framework to define the boundaries of epistemology. In the next section, I will elaborate what I take this framework to look like, which will help clarify the relationship I propose between the tasks of epistemology and the determination of its boundaries. After that, I will highlight some of the ways in which this approach to epistemology can produce significant benefits, both in terms of advancing our understanding of traditional epistemological problems, and in terms of opening up new areas of epistemological inquiry. In some cases of the latter, it is not so much a matter of

introducing *new* areas of inquiry as coming to understand why certain areas already being pursued are legitimately part of epistemology.

### **The Normative Structure of Epistemology**

I take the normative structure of any domain to be usefully characterized in terms of three elements.

- (1) The values that are relevant to the domain—in this case, “epistemic values,”
- (2) the set of kinds of objects that are evaluable in this domain, and
- (3) the set of value-transmitting relations.

I do not intend to defend any substantive claims about how each of these elements should be substantively filled out. My purpose is to explain how, by giving some specification of each of these elements, one can define some boundaries for epistemology. The point of such boundaries, unlike most real-world boundaries, is not so much to keep unwanted members out, but rather to discover the under-appreciated variety of flora and fauna within.

#### *Epistemic Values*

As one might expect, an understanding of epistemology along the lines elucidated above must begin with a specification of the values that we take to be central to it. What I mean by this is that we must articulate the values with respect to which we shall be evaluating objects from the “epistemic point of view.” (Such a list may certainly be provisional, and in line with my earlier endorsement of the methodology of reflective equilibrium, I think that any such list *should* be provisional, subject to revision as necessary to be consistent with changes made elsewhere in the structure. But this sort of commitment is not a necessary element of value-

driven epistemology, per se. For an insightful and detailed articulation and illustration of this method at work, see Catherine Elgin's (199?).)

Despite the centrality of such a list of values to the pursuit of value-driven epistemology, I have little to say here about what, precisely, those values are. I have argued for certain proposals along these lines elsewhere, as have others.<sup>2</sup> The standard view among epistemologists until fairly recently was that truth was the single, or at least singularly important, epistemic value. As philosophers tried to be more specific about what this actually means, all kinds of problems arose. It turns out to be very difficult to specify carefully the precise value represented by the innocuous seeming term "truth."<sup>3</sup> Furthermore, issues arise regarding the relative importance of gaining truths versus avoiding believing falsehoods, and how this balance should be determined.<sup>4</sup> And Catherine Elgin has recently gone so far as to argue that truth—strict, literal truth, at least—is not always as valuable as most have assumed.<sup>5</sup> Obviously, resolving these problems is of central importance to value-driven epistemology, at least insofar as one takes truth to be one of the main epistemic values. Since it remains fairly uncontroversial that truth is at least among the central epistemic values, these issues are front and center for value-driven epistemology.

But the dogma that epistemology has only truth as its central value has been challenged lately. Call this view Value Monism. Framing this issue is difficult, because one can think that there are valuable cognitive states or phenomena that are not valuable in the first instance because they are true, and yet still be committed to Value Monism. For example, a justified but false belief might be taken to be epistemically valuable to some degree, despite the fact of its

---

<sup>2</sup> See Riggs (200?), Kvanvig (200?)

<sup>3</sup> For a representative sampling, see DePaul, David, etc.

<sup>4</sup> Riggs

<sup>5</sup> Elgin, "True Enough,"...

falsehood. But a Value Monist would argue that the value of the justified false comes nonetheless from a kind of relationship to truth. Perhaps it is valuable because the person who holds the belief is *aiming* at the truth. Or perhaps the belief is supported by *truth-conducive* grounds. Whatever her theory of justification and her theory of the value of justification, a Value Monist will attempt to tie the value of justified false belief back to the value of truth, one way or another. So even someone who agrees that there properties of beliefs or persons that are of epistemic value might still maintain value monism by claiming that ultimately all such value is derived from the epistemic value of truth. This debate between value monism and value pluralism is also central to value-driven epistemology.

But my goal here is not to argue for a specific list of epistemic values, but rather to articulate the way such a list helps define and inform one's conception of epistemology as a whole. Philosophers often use the terms "goal" and "value" more or less interchangeably, and this allows for a certain ease of expression that would be hard to achieve otherwise. I will also talk about "epistemic goals" and epistemic values," but I want to make it clear how I take these to be related. I will take our epistemic values to comprise those properties which, when had by the appropriate object, make those objects epistemically good, *ceteris paribus*. Thus, truth is an epistemic value because, when a believed proposition has this property, it is epistemically good. Insofar as believing a true proposition is epistemically good, it is an appropriate epistemic goal.

If this all smells a little tautologous, bear in mind that I am not trying to do any heavy lifting here. I am simply clarifying my use of terms. It is not, perhaps, very informative to hear that Fness is an epistemic value iff something that Fs is an epistemically good thing, *ceteris paribus*. But it allows me to speak of an epistemic goal as a state of affairs in which some appropriate object has an epistemic value as a property. Besides which, I think we have a certain

kind of purchase on whether something is epistemically good that we don't necessarily have on the question of what our epistemic values are. Thus, we can use the empty-seeming definition to discover what our own commitments are in terms of epistemic values. For example, I think that we can intuitively rank different possible states of affairs in terms of their apparent or intuitive epistemic goodness. By choosing appropriate pairs of such states of affairs, we can "read off" the differences between the states of affairs that make a value difference, and possibly discover that certain properties are epistemically valuable that we had not considered before.

Let me give a brief example of what I mean by the last comment. My own way of determining what I take to be the correct set of epistemic values is by considering what I would include in the description of the best possible epistemic life. This would presumably include, at the least, believing all true propositions and believing no false propositions. Beyond this, no doubt opinions will vary considerably. On my own view of knowledge, it would be better to not only believe all truths, but also to know all of them that are knowable. So if we imagine being in such a state of knowing all knowable truths and having no false beliefs, can we imagine a state that is any better epistemically? One's answer to this will depend on a number of factors, but among these will be one's account of understanding. It is fairly uncontroversial, I think, that a kind of systematic explanatory coherence among one's beliefs is an epistemically desirable thing. If it is possible to know all knowable truths and fail to have this kind of systematic coherence among those known truths, then it would seem that something of epistemic value is still missing. And, if one takes understanding to consist in this kind of systematic or explanatory coherence, then one can conclude that understanding is an additional epistemic value beyond both truth and knowing.

But some might argue that such coherence or understanding consists primarily or wholly in *beliefs* about the interrelations among one's beliefs. If so, then anyone who knows all knowable truths (assuming, plausibly, that such things are knowable) will also know these interrelations. Thus, the value that was claimed to be missing in that scenario is present after all. This amounts to reducing the property of coherence to a set of true (meta-)beliefs, and thus the value of such coherence to the value of believing or knowing those particular truths. This strategy would have obvious appeal to a value monist as described above.

There are many ways this debate might play out, but that is not of concern to the present topic. I simply wanted to show how a consideration of the comparative intuitive epistemic goodness of two possible states of affairs could suggest that something is an epistemic value that was not necessarily on the list to begin with.

So there are plenty of issues that need to be resolved in order to develop a defensible list of epistemic values. Assuming that truth is an epistemic goal, how precisely do we formulate the goal that this implies for us? Just how valuable is truth anyway? How are the value of truth and the disvalue of falsity to be balanced? Is value monism true or are there other properties of genuine epistemic value that do not derive their value from the value of truth? All of these are of central importance to value-driven epistemology, and all of them have implications for epistemology broadly construed, as well as for the theory of knowledge.

### *Objects*

The second element to fill out in the normative structure of epistemology is the set of kinds of objects that are appropriately evaluable in that domain. This is, in some ways, the least philosophically interesting element in the structure being defined here because in one sense it hardly matters what we put in here. After all, one may evaluate *anything* with respect to a given

set of values. Many objects will simply be neutral with respect to those values. For example, one might ask how good a stone is with respect to the value of truth. Since there is no interesting way in which a stone can be related to the anything bearing the property of truth, the stone is neither good nor bad with respect to truth. But there is nothing stopping us from asking the question.

On the other hand, it is worth pointing out that there are different possible objects of evaluation within epistemology in particular. Epistemology in recent decades, especially perhaps in the immediate aftermath of Edmund Gettier's famous paper<sup>6</sup>, has focused very narrowly on the notions of "knowledge" and "justification." Furthermore, the canonical characterization of knowledge during that time was that knowledge=justified true belief—the 'JTB' account of knowledge. This has the perhaps unintended effect of focusing our attention on a particular object—the belief. To say that knowledge is justified true belief is to suggest that instances of knowledge are simply instances of true beliefs that have a further property—"justification." We talk, after all, of justified *belief*, and thus reinforce the idea that the only object being evaluated here is a belief.

Interestingly, though this terminology has persisted even until now, actual analyses of knowledge and justification have long cited properties of things other than beliefs. So, epistemologists have implicitly acknowledged that S's knowing that p may require that objects other than the belief itself must have certain properties. Examples include the following:

- The process that produces p in S must be reliable
- S's cognitive behavior relevant to coming to believe that p must have been epistemically blameless
- S must have come to believe p in an epistemically responsible way
- S must base her belief that p on adequate grounds
- S must have intellectual virtue V

---

<sup>6</sup> Gettier, ""

Each of these is an example of a commonly cited necessary condition for either knowledge or justification. Notice that in each case, the required property is not predicated of a belief, but of a belief-producing process, S's behavior, S's intellectual character, and so on. It is true that we could concoct a property of a belief in each case that seems to capture the same idea. For example, epistemologists often talk about "reliably-produced belief," as though being "reliably-produced" was a property of the belief. And, of course, it is a property of the belief. But this point obscures more than it illuminates, especially within the domain of value-driven epistemology. Does the fact that S has come to believe that p in an epistemically responsible way make the belief that epistemically better? *Something* in the neighborhood seems epistemically good, but is it the *belief*? In order to answer that question, we have ask how the value inherent in the properties listed in our set of epistemic values gets transmitted to the belief. Discussion of that process must wait until the next section. But for now, the important thing to note is that once we acknowledge that meeting the conditions of knowledge might require that we evaluate objects other than beliefs, we open the door for new source of value that knowledge might have. Since figuring out whether, to what extent, and precisely how knowledge is valuable is a question that presumably is important to epistemology even narrowly construed, this is a significant point.

Another important implication of the realization that many objects are at least potentially evaluable from "the epistemic point of view" is that it widens the scope of epistemology. Given the contemporary focus on knowledge and justification, and the further narrowing of looking primarily at beliefs, it is hard to justify the pursuit of other epistemic game. Consider how the recent rise of so-called "virtue epistemology" has gone. While the views that fall loosely under that heading do not form a homogenous group, one thing they tend to have in common is that

they are virtue theories *of knowledge*. They all put what they each call “virtues” at the forefront of theoretical consideration, but only as part of a theory of knowledge. Few are those who are willing to say that epistemic virtues are worthy of theoretical analysis whether or not they are part of a successful theory of knowledge.<sup>7</sup>

Similarly, once one takes seriously the idea that sets of beliefs are possible objects of epistemic evaluation *in their own right*, then questions about coherence, understanding and the like are also seen to be of theoretical interest independent of any contribution they might make to a theory of knowledge or justification. There are, in fact, plenty of objects that we can evaluate from the epistemic point of view that might have little or nothing to do with knowledge or justification. Defining epistemology more broadly than the “theory of knowledge” as the structure articulated here does both highlights this general point and also helps point the way to where there is interesting epistemological work to be done.

#### *Value Transmission*

I take it to be a very general question how any epistemically valuable state, phenomenon, or what have you derives its value from our fundamental epistemic goals. But to make the discussion less abstract, let us consider for now the transmission of value from our fundamental epistemic values to supposedly epistemically valuable state of knowledge. For simplicity, we can even assume for now that the only fundamental epistemic values are having true beliefs and avoiding believing falsehoods. For our purposes, we can treat these epistemic goals as if they were intrinsically valuable. I am not claiming that they really are intrinsically valuable, but I am merely pointing out that I am interested in the value chain that begins with those values and goes on to infuse the state of knowledge with epistemic value. Thus, the source of the value inherent in those goals is not of concern to this project.

---

<sup>7</sup> Exceptions include Kvanvig, Riggs, ??

But how is the value of those ends “transmitted” to the state of knowledge? Presumably, there must be a chain of relations between the epistemic goals and the state of knowledge, each of which transmits some or all of the value inherent in our ultimate epistemic goals to the next object in the chain. But what sorts of relations are these? Not every relation between a valuable thing and something else transmits the value of the one to the other. The relation of “being to the side of” will not usually render an object X valuable even when it bears that relation to something valuable. Value that comes to an object via its relation to another object is usually called “extrinsic value.” So, one way to answer the question before us is to ask what the different kinds of extrinsic value are. While I do not claim to give an exhaustive list, I will describe the different kinds of extrinsic value that I think are involved in the value chains I am interested in.

Compared to the literature on intrinsic value, there is very little discussion of extrinsic value. Philosophers often employ the notion of extrinsic value, but they rarely have much to say about the nature and kinds of extrinsic value. Therefore, there is neither a standard taxonomy nor a standard terminology to which to appeal. I shall borrow both from recent discussions of extrinsic value.

The most familiar type of extrinsic value is *instrumental* value. In fact, extrinsic value is often equated, wrongly, with this type of value. Instrumental value is value that something has because it is productive of, or is a means to some further value. For instance, money has instrumental value because it can be exchanged for other things we value. Typically, money has no other value than this, though one could I suppose appreciate the beauty of a particular coin or bill (speaking strictly of non-American currency, obviously).

Another familiar type of value is what I will call *teleological* value. This is value that something has by virtue of being “aimed at” something else of value. Though virtually anything

can have instrumental value, teleological value is limited to objects that are capable of intentionality. The bearer of teleological value must be something which has “aims.” This can be interpreted in any of several ways. We can think of aims as desires, goals, ends or purposes, but the important point is that the bearer of teleological value be intentionally directed toward something that is itself valuable. So, for instance, a person’s motive for action X will have teleological value if the end or goal of the motive is itself of positive value. This will be true whether or not action X succeeds in achieving the end or goal of the motive. Suppose that Jane buys lunch for a homeless person from the motive of benevolence, the end or goal of which is, let us say, to relieve suffering. It may happen that the homeless person has an ulcer which is aggravated by the food Jane provides him, so that he is actually worse off than if she had let him remain hungry. While this may affect the moral value of the action, depending on one’s moral theory, it does not affect the value of the motive. Insofar as Jane was acting so as to relieve suffering, her motive was valuable even if it does not produce the good at which it was aiming. That is because intentional states in general derive value from that at which they are aiming. That is what I call teleological value.

An object can also be valuable if it is a necessary part of some greater whole that has value. Call this *contributory* value. If a MRI machine is valuable because it helps save lives, then the individual components of it that make it work are also valuable. It is important that the part in question be necessary. If, for example, a particular MRI machine has a post-it note reminding the technician to pick up her dry-cleaning this afternoon on it, the post-it note does not thereby derive value by being part of the machine. This opens thorny questions about how to determine necessary vs. unnecessary parts of a functioning whole, none of which am I going to address

here. However such issues are sorted out, the parts that are necessary to a valuable whole will thereby derive contributory value from the whole.

And finally, there is what I will call *indicative* value. A common example of this is also from the medical field. Suppose that my doctor orders an X-ray of my lungs because she thinks that I might have a tumor. The X-ray comes back, she examines it, and determines that I do not, in fact, have a tumor in my lungs. It would be perfectly natural to say in this context that my X-ray was “good.” Thus, the X-ray derives value from the value of my not having a tumor.

In all these cases, the kind of value is identified by the relation to value through which the value is derived. An object has instrumental value by virtue of its instrumental relation to some other thing of value. Similarly with teleological, contributory, and indicative value. The kind of value is characterized in terms of the relation through which value passes from some other object of value to the object with the kind of value being discussed.

Having distinguished these different kinds of value, or alternately, these different relations through which extrinsic value is transmitted, how do we account for the value that knowledge has? As an example of value-driven epistemology at work, I will try to articulate the value chain of a simplified version of reliabilism. This will both offer an example of the kind of axiological investigation that can help guide our theory of knowledge, and it will also provide a substantive result, albeit one that has been shown elsewhere (Zagzebski, Kvanvig, Jones, Riggs, etc., get refs)—namely, that simple reliabilism fails to account for the additional value that knowledge has over mere true belief.

We must keep in mind throughout this discussion that the task is not merely to find some value that knowledge holds for us, but rather to account for the value knowledge holds above and beyond the value of having a true belief. Since knowledge is nearly universally taken to entail

true belief, having knowledge always provides one with the value of having a true belief. The question is what value is added to true belief by virtue of its being an instance of *knowledge* rather than a true belief that falls short of knowledge in some way.

According to simple process reliabilism (SPR), a belief is justified iff it is the product of a reliable process. If the belief is also true, then the belief is an instance of knowledge. As we already mentioned, the reliably-produced true belief will be epistemically valuable in virtue of its being true. Thus, the question remains, what additional value is had by a true belief that is reliably-produced? Since we are still operating under the assumption that our only epistemic goals are having true beliefs and avoiding false beliefs, the value a belief gains by virtue of being reliably-produced must be traced back to one or both of these goals.

Reliabilists tend not to talk very much about the value of knowledge,<sup>8</sup> so we are left to reconstruct how the value chain is supposed to go on their behalf. Intuitively, the idea seems to be that a justified belief is more valuable than an unjustified belief precisely because it is the product of a reliable process. We want all our beliefs to be produced by a reliable process, so the ones that are so produced are more valuable than the ones that are not. As we shall see, this intuitive picture breaks down when we bother to look at the details of how this is supposed to work.

To begin with, it is fairly easy to see how a particular belief-forming process can be epistemically valuable. Insofar as a process tends to produce true beliefs rather than false ones, to that extent it is instrumentally good with respect to our epistemic ends. Thus, a process can derive value from our epistemic ends by virtue of being instrumentally related to them. But notice that only process-*types* can derive such value. Only process-*types* can be meaningfully said to *tend to* produce true beliefs rather than false ones. Process tokens necessarily produce a

---

<sup>8</sup> Greco and Sosa are exceptions to the rule

single belief at a particular time. This belief will be either true or false. There is no sense on which a token process can be said to “tend to” produce true rather than false beliefs. A token process that produces a true belief derives epistemic value instrumentally thereby, and a token process that produces a false belief derives epistemic disvalue thereby.

So a process-type is epistemically valuable insofar as some percentage of its tokens (both actual and possible, but more on this later) produce true beliefs rather than false ones. A type is valuable insofar as the tokens that comprise the type are valuable, so a process-type derives value from its tokens by simple aggregation. If most of the tokens are valuable, then the type is valuable as well.

Thus far, the value-derivation picture looks like this.

$r_1 \rightarrow$  true belief, so is epistemically valuable  
 $r_2 \rightarrow$  true belief, so is epistemically valuable  
 .  
 .  
 .  
 $r_3 \rightarrow$  false belief, so is epistemically disvaluable  
 .  
 .  
 .  
 $r_N \dots$

$r_1 \dots r_N$  are token belief-forming processes  
  
 $R$  is the type, which is the collection of all the belief-forming process tokens,  $r_1 \dots r_N$ .  
  
 The value of  $R$  is simply the sum of all the values of  $r_1 \dots r_N$ .  
  
 Or more formally,  

$$V(R) = V(r_i \sum_{i=1}^N)$$

We have shown how a belief-forming process-type can derive value from our epistemic ends, but now we must complete the value derivation by showing how this value makes its way

down to a particular belief that is produced by a token of such a type. Because we haven't shown that knowledge is valuable until we can show that any actual belief that meets the conditions of knowledge has the value in question.

But this is where the value-transmission picture gets murky. Given that a process-type R is valuable, and even given that belief p is produced by a token-process r of that type, how does the value of R get transmitted to p? What relationship is there between p and R that could transmit the value of R to p? The only relationship they bear to each other at all is that p is the product of a token of type R. There is no direct relationship between them. Their only relationship is mediated through the process token that actually produces p. That suggests that any relation through which value could be transferred from R to p must go through r. But now the circularity of this account becomes apparent. The value of R comes entirely from the aggregate value of  $r_1$  through  $r_N$ . That value cannot then be "turned around," as it were and then be transmitted back to the process-tokens, thereby increasing their value. If the type is valuable because it has mostly valuable tokens, you cannot then say that the tokens are more valuable because they are members of that type. The value would be going from the tokens to the type and back again, with no additional value added at any point. Thus, a process-token that produces a true belief cannot be made more valuable by virtue of the fact that it is an instance of a valuable process-type. Simple process reliabilism, then, cannot account for the transmission of value from our epistemic ends of having true beliefs and avoiding falsehoods.

It may well be that more sophisticated versions of reliabilism are more successful at accounting for the full value of knowledge. This illustration of tracing the value chain from the item of alleged value (knowledge) to that from which it allegedly derives such value (truth) is meant only to show how value-driven epistemology can contribute to the traditional theory of

knowledge. It places an additional constraint on such theories—that they must be able to show that knowledge is as valuable as we take it to be pre-theoretically. It also gives us the tools to determine whether a theory meets that constraint or not.

## **Conclusion**

Having spelled out the normative structure of epistemology, we are now in a position to answer the question posed early on in this essay: “What is Value-Driven Epistemology?” As I see it, value-driven epistemology is both a way of construing epistemology as well as a methodology for pursuing it. First, value-driven epistemology accepts that the bounds of epistemology are given by the normative structure described in the last section. Thus, any property, phenomenon, state or object that has a place in that structure is an appropriate subject of epistemological investigation. Second, a preferred way of beginning such an epistemological investigation is by asking what epistemic values are involved, and how those values get transmitted or transferred to the object under investigation. This will by no means exhaust the epistemological work to be done, but it will both provide important details about the phenomenon being investigated as well as placing additional constraints on the possible answers to other questions that might be asked. The illustration of the value-derivation for Simple Process Reliabilism is one example of this.

My main aim in this paper has been to describe the value-driven approach to epistemology, rather than to defend it. A lot of contemporary, and even historical, epistemology could reasonably be categorized as value-driven, but I think it is helpful to make this way of doing epistemology explicit, the better to pursue it consistently, and to see how it relates to other approaches to epistemology. Nevertheless, I hope that along the way some of the illustrations

and examples have indicated the advantages of doing epistemology this way, at least as one way among others. In this conclusion, I want to simply highlight a few of the ways I think the value-driven approach can and has contributed to epistemology generally.

To my mind, the most significant contribution that value-driven epistemology has to make is that it keeps epistemological theorizing on the track of those things that really matter to us. For example, Timothy Williamson and Jonathan Kvanvig have each argued recently<sup>9</sup> that post-Gettier theories of knowledge have produced some incredibly complex and arcane criteria for knowledge, and that it is far from clear that we do or should value meeting such criteria. And yet, if knowing is necessarily something of value to us, then such analyses must be misguided.

[T]he difference between knowing and not knowing is very important to us. Even unsophisticated curiosity is a desire to *know*. This importance would be hard to understand if the concept *knows* were the more or less ad hoc sprawl that analyses have had to become; why should we care so much about *that*? (Williamson 2000, pp. 30-31)

Kvanvig goes so far as to argue that this mismatch between the “ad hoc sprawl” of post-Gettier criteria of knowledge and something we care about might indicate that there simply is no account of knowledge that both defeats Gettier and is clearly more valuable than some other state that falls short of knowledge.

On the one hand, the variety of cases in which one can fail to know and yet have a justified true belief inclines one toward more complex, ad hoc, and gerrymandered proposals. On the other hand, the felt need to address the question of the value of knowledge over its subparts leads one toward simpler proposals in which the value of knowledge is intuitively obvious. The twin desiderata on a theory of knowledge, the desiderata of accounting for both the nature and value of knowledge, threaten in this way to become the Scylla and Charybdis that sink the project entirely.

Whether Kvanvig is right or not about the ultimate success of any theory of knowledge, his demand that any such theory must meet the desideratum of accounting for the value of

---

<sup>9</sup> (Williamson 2000), pp. 30-31; (Kvanvig 2003), pp. ???

knowledge keeps such theories from being driven by the need to avoid counterexamples into territory that we no longer recognize as desirable.

The increased attention to issues surrounding epistemic value in recent years has, I think, left its mark on the theoretical territory. In the development of so-called “virtue epistemology” especially, explicit appeals to and debate about epistemic values has been influential. Linda Zagzebski, in her highly regarded book *Virtues of the Mind* (Zagzebski 1996), holds herself accountable to the demand that her criteria of knowledge must preserve the value such a state is supposed to have over and above mere true belief, and this demand places non-trivial requirements on those criteria. Furthermore, since for Zagzebski the very virtues that are at the center of her account are defined in part by what the central epistemic value is, this issue is crucial to the entire enterprise of her epistemology.

Other virtue epistemologists, like John Greco and Ernest Sosa, have also focused more and more on the nature and importance of epistemic values as their theories have become ever more explicitly dependent on the notion of epistemic virtue. And, more importantly, each of them has, like Zagzebski, defended his view by arguing that his account of knowledge can account for its special value. By taking this constraint seriously, they have both acknowledged the important role that epistemic values have to play in epistemology, and they are using value-driven types of reasoning to defend their views.

But perhaps the greatest potential effect of value-driven epistemology is its openness to new epistemological investigations that are not tied to accounts of either knowledge or epistemic justification. As important as those concepts are to epistemology, they do not exhaust the range of epistemic evaluations that are worthy of study. A broad conception of epistemology, like the one implicit in a value-driven approach to epistemology, legitimates investigation of these other

concepts. Furthermore, once one realizes that knowledge and justified belief, while epistemically valuable, need not be the *only* epistemically valuable things around, one has a reason to actively look for other epistemically valuable states. This should lead to a much richer picture of what constitutes the epistemic good life.

The best illustration of this process that I know of is Kvanvig's work on knowledge and understanding<sup>10</sup>. Kvanvig defends the strong thesis that knowledge is actually not more valuable than various combinations of its subparts. While that is an interesting debate in its own right, Kvanvig goes on to use this "value lacuna" to suggest that there is perhaps a different epistemic state that is valuable in much the way we take knowledge to be valuable. The state he has in mind is that of understanding. For whatever reason, this obviously valuable and obviously epistemic notion has received very little philosophical attention, especially in recent decades. I take the narrow, problem-oriented approach to epistemology to be largely to blame for this, along with the assumption that knowledge is the paradigm example of epistemic excellence. But Kvanvig's work challenges this view in two ways. First, he raises serious questions about whether knowledge really is the paradigm of epistemic excellence. If having knowledge is not all that valuable after all, then achieving epistemic excellence may require more than, or at least something other than, the accumulation of knowledge. And second, Kvanvig puts the need to find the bearers of epistemic value at the fore of his epistemological theorizing. When he concludes that knowledge does not bear that value, he looks elsewhere within his broad conception of epistemology to find it. This is value-driven epistemology at its purest.

---

<sup>10</sup> See (Kvanvig 2003)

## Works Cited

- Alston, William. (1989) "Concepts of Epistemic Justification," in *Epistemic Justification*, William Alston (Ithaca, NY: Cornell University Press).
- Alston, William. (2005) *Beyond Justification* (Ithaca, NY: Cornell University Press).
- Chisholm, Roderick. (1989) *Theory of Knowledge*, 3<sup>rd</sup> edition (Englewood Cliffs, NJ: Prentice Hall).
- David, Marian. (2001) "Truth as the Epistemic Goal," in *Knowledge, Truth, and Duty*, Matthias Steup, ed. (Oxford: Oxford University Press).
- DePaul, Michael. (2001) "Value Monism in Epistemology," in *Knowledge, Truth, and Duty*, Matthias Steup, ed. (Oxford: Oxford University Press).
- DeRose, Keith. (1999) "Contextualism: An Explanation and Defense," in *The Blackwell Guide to Epistemology*, John Greco & Ernest Sosa, eds. (Oxford: Blackwell).
- Elgin, Catherine. (1996) *Considered Judgment* (Princeton: Princeton University Press).
- Elgin, Catherine. (2005) "True Enough"
- Gettier, Edmund. (1963) "Is Justified True Belief Knowledge?" *Analysis* 23.
- Greco, John. (2003) "Knowledge as Credit for True Belief," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Michael DePaul and Linda Zagzebski, eds. (Oxford: Oxford University Press).
- Jones, Ward. (1997) "Why Do We Value Knowledge?" *American Philosophical Quarterly* 34.
- Kvanvig, Jonathan. (1992) *The Intellectual Virtues and the Life of the Mind* (Savage, MD: Rowman & Littlefield).
- Kvanvig, Jonathan. (1998) "Why Should Inquiring Minds Want to Know?" *The Monist* 81.
- Kvanvig, Jonathan. (2004) *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press).
- Riggs, Wayne. (2002) "Beyond Truth and Falsehood: The Real Value of Knowing that P," *Philosophical Studies* 107 (1).
- Riggs, Wayne. (2003) "Balancing our Epistemic Goals," *Noûs* 37 (2).

Riggs, Wayne. (2003) "Understanding Virtue and the Virtue of Understanding," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Michael DePaul & Linda Zagzebski, eds. (Oxford: Oxford University Press).

Rorty, Richard. (1967) "Preface" to *The Linguistic Turn*, Richard Rorty, ed. (Chicago: University of Chicago Press).

Zagzebski, Linda. (1996) *Virtues of the Mind* (Cambridge: Cambridge University Press).