2017-11-20

Reconsidering Epistemic Evaluation

Lance Aschliman
University of Miami, lkaschliman@yahoo.com

UNIVERSITY OF MIAMI

RECONSIDERING EPISTEMIC EVALUATION

By

Lance Kelly Aschliman

A DISSERTATION

Submitted to the Faculty
of the University of Miami
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Coral Gables, Florida

December 2017

UNIVERSITY OF MIAMI

A dissertation submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy

RECONSIDERING EPISTEMIC EVALUATION

Lance Kelly Aschliman

Approved:

_____
Michael Slote, Ph.D.
UST Professor of Ethics

_____
Harvey Siegel, Ed.D.
Professor of Philosophy

_____
Berit Brogaard, Ph.D.
Professor of Philosophy

_____
Elijah Chudnoff, Ph.D.
Associate Professor of
Philosophy

_____
Allan Hazlett, Ph.D.
Professor of Philosophy
Washington University in St. Louis

_____
Guillermo Prado, Ph.D.
Dean of the Graduate School

ASCHLIMAN, LANCE KELLY   (Ph.D., Philosophy)
<u>Reconsidering Epistemic Evaluation.</u>   (December 2017)

This project is an examination of epistemic evaluation, of what it is and how we should think about it in our overall normative theories. The overarching goal of this dissertation is to argue that it is theoretically fruitful to frame epistemic evaluation in a specific way, as a kind of "critical domain" or evaluative practice that takes the constitutive goals of inquiry to be fundamentally valuable. In due course, I argue that *knowledge* and not mere true belief is fundamentally valuable from the epistemic perspective and that, when we see epistemic evaluation in this way, a number of difficult theoretical pieces fall into place, including answers to questions concerning the so-called "value problem of knowledge" and to questions raised about the value of epistemic justification. Some may object to seeing epistemic evaluation as a critical domain on the basis of considerations arising out of the ethics of belief, considerations concerning what we all-things-considered *ought* to believe. I provide a framework for thinking about these issues and argue that, at the end of the day, any worries about the adequacy of the critical domain picture with respect to the ethics of belief are ill-founded and that the view presented is a live theoretical option.

Table of Contents

Introduction

Epistemologists engage in evaluation. They evaluate individual beliefs as to whether they constitute knowledge, whether they are justified, whether they are supported by evidence, whether they have been formed via reliable cognitive processing, whether they are "warranted," etc. They evaluate sets of beliefs as to whether they are consistent, whether they "cohere," whether they constitute an understanding of a particular subject, whether they are explanatorily fecund, etc. They evaluate forms of reasoning as to whether they have conformed to epistemic rules. They evaluate persons or agents as to whether they act in epistemically virtuous ways or whether they have habits or dispositions that are epistemically virtuous. Evaluation abounds in epistemology. Surely enough, epistemologists do more than evaluate; they also engage in descriptive projects. An epistemologist might attempt to uncover the *nature* of knowledge, or the *nature* of the basing relation, or the *nature* of *a priori* cognition. But, even these descriptive research projects are conducted against a background of an evaluative practice which assumes that certain things can be better or worse from a particularly epistemological perspective.

Given this, we can ask the following questions: What kinds of evaluations are distinctively epistemic in character? When we say that, for example, epistemologists evaluate beliefs as to whether they constitute knowledge, why is the positive evaluative status of constituting knowledge a particularly *epistemic* one? One might be tempted to answer questions like these by appealing to the etymology of the word "epistemology" as being related to a theory of knowledge, but this cannot be the full story given the wide

range of the epistemologist's objects of study. The real question is the following: What distinguishes the many distinctively epistemic evaluative considerations from other sorts of evaluative considerations, be they moral, prudential, aesthetic, or otherwise? Asking such a question is not to imply that epistemic considerations are never also relevant to moral or prudential or aesthetic ones. After all, it often does seem that having knowledge of certain things does contribute to, say, living the good life. Presumably, though, we cannot appeal to *that fact* in explaining why evaluations regarding knowledge are particularly epistemic. Some bigger-picture story needs to be told.

In the first chapter of this dissertation, I will discuss an initial prospect for this bigger-picture story. I will suggest that the practice of epistemic evaluation is best seen as what Ernest Sosa (2007) calls a "critical domain." One major motivation for the project undertaken here is that Sosa's idea of epistemology as a critical domain has some severely underappreciated implications, even in what Sosa himself draws from the idea. I will do two things in Chapter 1. First, I will explain, in detail, what a critical domain is in a way that will be most useful for the theoretical purposes here. Second, I will suggest that epistemology is a critical domain structured by the constitutive goals of inquiry. The remaining chapters build on this view, situating it within a wide range of literature on both epistemic and non-epistemic forms of evaluation.

If the account proposed in Chapter 1 is the right one, not only will it help us answer the questions posed above about what makes certain evaluative considerations particularly epistemic, it will also help uncover some serious errors in past epistemological theorizing. For example, I will argue in Chapter 2 that it is a mistake to think, as many do, that the most fundamental value within epistemology is that of true

belief. I will argue that the view sketched in Chapter 1, where *knowledge* is taken to have a distinctive kind of fundamental epistemic value, puts us in a much better position to understand a version of the so-called "value problem of knowledge." In fact, I will suggest that mere true belief is best seen as having a very minimal kind of epistemic value. This will help make it clear what, exactly, we are after in our epistemic evaluative practice and what sorts of considerations make our epistemic evaluations important.

The remaining chapters focus in on the implications for a certain notion of epistemic justification. Though I argue that no special place can be accorded to mere true belief, I argue in Chapter 3 that justified belief does play a distinctive role within the epistemological critical domain. It will be shown that epistemic justification is valuable with respect to the goals of inquiry in virtue of exhibiting a kind of inquiry-relevant capacity. Chapter 4 then considers the kinds of connections that epistemic justification might have to more genuinely normative, guidance-giving prescriptions on belief. The ultimate aim of these chapters to make it clear why it is important for us, both in theory and practice, to see epistemology as a critical domain structured by the goals of inquiry.

Chapter 1: Epistemic Evaluation and the Goals of Inquiry

In this chapter, I suggest that epistemic evaluation is best seen as an instance of what Ernest Sosa (2007) calls a "critical domain." In the first section, I examine Sosa's explicit remarks about these domains and extend them in order to establish a more general and clearer picture of what critical domains are supposed to be. Most importantly, I make two distinctions: (1) the distinction between fundamental and derivative value in a domain and (2) the distinction between domain-internal value and genuine value. In the second section, I take this general formulation of critical domains and apply it to epistemology, arguing that epistemic evaluation is a specific kind of critical domain framed by the constitutive goals of inquiry. I do this in a way that remains neutral in the debate between the monists and pluralists about fundamental epistemic value. In the final section, I field a number of initial objections that might be brought out against such a view of epistemic evaluation, setting the stage for an extended discussion of the view in the chapters to come.

1.1: What is a Critical Domain?

What does it mean for some evaluative practice to be what Ernest Sosa calls a "critical domain?" In developing his own distinctive brand of virtue epistemology, Sosa describes a critical domain as a "set of interrelated entities evaluable through correspondingly interrelated values." (2007, pg. 73) Purportedly, critical domains admit "values of two sorts: the derivative, and the fundamental – that is to say, the derivative or fundamental *for that domain*." (pg. 70) Of course, by themselves, these quotes do not tell

us much about what Sosa has in mind. The goal of this section is to make his general notion of a critical domain clearer, specifically how certain kinds of things and values within them are "interrelated" in order to see how the idea is supposed to apply to epistemology later on.

When further explaining the concept of a critical domain, Sosa resorts to description by example and it will be easiest to follow him here as a jumping-off point:

> Consider the world of coffee – of its production, elaboration, and consumption. One central value organizes the critical assessment distinctive of that domain. I mean the value of liquid coffee that is delicious and aromatic. Think of the assessment of coffee beans, fields, coffee machines, baristas, ways of making liquid coffee, plantations, harvests, etc. What organizes all such evaluation, the value at the center of it all, from which the other relevant values are derivative, is the value of good coffee, of liquid coffee that is delicious and aromatic. (2007, pg. 73)

According to Sosa, there are many kinds of things within the "world of coffee" that are each uniquely evaluable, but are yet evaluable, in some sense, from the very same perspective. That evaluative perspective comes with certain fundamental values that ultimately structure all of the relevant evaluations in the domain. For the case of coffee, Sosa claims that the fundamental value of delicious, aromatic liquid coffee structures all relevant coffee-related evaluations.

The general interrelation among the entities (or possible entities) comprising a critical domain seems to be something like this: If a certain kind of thing like delicious, aromatic coffee is valuable from a certain perspective, then other kinds of things will also take on value through derivative value relations. Different kinds of things are "interrelated" in a domain by way of these value relations. For example, since delicious, aromatic coffee is good from a certain perspective, then certain kinds of baristas are also good (e.g., those who engage in producing such coffee). Similarly, since delicious,

aromatic coffee is valuable from a certain perspective, then certain kinds of bean harvests are also valuable (e.g., those in which the beans harvested might someday contribute to producing such coffee). Ultimately, within a critical domain, the value accruing to certain fundamentally valuable things *grounds* the value of other related entities, explaining the value of the more derivative, yet related, items.

As the examples above suggest, derivative value relations include (at the very least) instrumental or contributory relations where entities are valuable as a means to achieving other, more fundamentally valuable things. I will have more to say about the nature and scope of these derivative value relations in the coming chapters but the idea here is that the values generated by such relations are grounded by the most fundamental ones. As Tamar Szabó Gendler (2002) would say, when two entities enter into certain kinds of relations with one another (e.g., a particular bean harvest and the delicious, aromatic coffee it produces), one may asymmetrically "borrow lustre" from the other. By entering into these relations, things become good *in virtue of* the more fundamental goodness of the other kinds of things. These derivative value relations are present in any critical domain, whether it be the realm of coffee, competitive tennis, epistemology, or otherwise.

So, derivative values within a domain, at bottom, are grounded in the fundamental ones. What, though, about those fundamental values? Are they grounded in any independent, value-theoretic facts? Why, for example, is delicious, aromatic coffee fundamentally valuable from the perspective of "the world of coffee" and not something else? Is that even a coherent question? One thing that does seem to be clear is that it is our own interests and goals that explain our desire to participate in particular evaluative

practices with their associated fundamental values. For the case of coffee, it is because some of us have an interest in delicious, aromatic coffee and engage in the goal-directed activity of consuming delicious, aromatic coffee that we are interested in a critical domain that takes delicious, aromatic coffee as fundamentally valuable. Yet, we need to be careful that we do not conflate the etiological question about what *causes* us to evaluate things in a certain way and the more metaphysical question about what grounds the most fundamental values in a domain. Our own interests and goals might play an important role in answering the first kind of question while not playing any role in answering the second.

I am inclined to think that the most fundamentally valuable kinds of things in a domain take on their valuable status *constitutively* of the domain itself, perhaps even being fundamentally valuable from that perspective as a matter of conceptual truth. There may not be anything metaphysically interesting to say about why, for example, delicious, aromatic coffee is fundamentally valuable within the "world of coffee." It just *is*. Subsequently, once certain fundamental values within a domain are picked out, it is possible for us to tell a further story about whether or not such a practice is interesting or relevant to us (i.e., whether we have some relevant interest or goal with respect to those fundamentally valuable things). Our interests and goals are what will ultimately make those evaluative practices things that we are motivated to participate in. This is all consistent with the existence of nearly any number of critical domains, including very strange domains that would, say, take bland, unsavory coffee to be fundamentally valuable. This should not be seen as a problem, though. Again, even if a domain exists

and constitutively has the fundamental values it has, it is a further question as to whether or not we should find such a domain interesting.

Take another example of a critical domain: the "world of competitive tennis." That world is structured by the fundamental value of winning tennis matches.[1] It is constitutive of that domain that winning tennis matches is its fundamental value. Why should people care about such an evaluative practice? Winning tennis matches is often the end of an associated goal-directed activity, something that people who participate in that evaluative practice have an interest in. Of course, saying that winning tennis matches is fundamentally valuable within the "world of competitive tennis" is not to say that victorious matches are the only good things from that perspective. Many other things become derivatively good in that "world" in virtue of those fundamentally valuable things. Certain tennis rackets are good from that perspective, certain shots and styles of play are good (given the contingent features of a situation), certain attitudes are good, etc.

Now, by themselves, fundamental and derivative values cannot give us a complete and general picture of these domains. We need to understand that evaluations within critical domains are structured by more than positive values; they are also structured by *disvalues*. That is, sometimes things are genuinely *bad* from a certain perspective or within a "world" and, when this is the case, the badness of those things is a property over and above the mere lack of goodness. Consider the coffee example again. When someone drinks a really bad cup of coffee, it is not merely the case that the coffee

---

[1] This is, of course, to oversimplify the domain. Plausibly, to win matches by disqualification would not be to satisfy the most fundamental goal. This is to say that the most fundamental goal would be something more like winning *in a certain kind of way* (perhaps by playing matches to completion, within the rules, in an elegant way, etc.).

that is being consumed is evaluatively neutral and merely lacks the property of being good. It is a genuinely bad cup of coffee! We cannot derive this result from a picture of critical domains that lacks the integration of negative valuation. Just as there are fundamental and derivative *values* within critical domains, there are also fundamental and derivative *disvalues*.

We explain the importance of a domain through its relations to certain interests or goals that are directed at its most fundamentally valuable things. Something similar can be said of a domain's most fundamental disvalues. For example, it is not usually the case that people who are interested in drinking delicious, aromatic coffee are *only* interested in that goal-directed activity. They are usually correspondingly interested in *avoiding* coffee that tastes a certain way. Consider the following scenario: Machine *A* and Machine *B* are just as likely to produce coffee that is delicious and aromatic (suppose that both produce good cups of coffee about ½ of the time). When Machine *A* does not produce good coffee, it produces coffee that is merely evaluatively neutral, mostly tasteless but unoffensive. Machine *B*, however, occasionally produces coffee with very pungent off-flavors. Now, if there were no genuine negative valuation within the "world of coffee," a coffee connoisseur would have no taste or aroma-related reasons to prefer a cup from Machine *A* instead one from Machine *B*. Yet, evaluators of coffee are and should be risk averse in this way. In the same way that a possible evaluator's positive interests can explain the importance of the fundamental values of the critical domain, the associated negative interests explain the importance of these fundamental disvalues.

The general picture is this: Critical domains are evaluative practices structured by a certain set of fundamental values and disvalues. For each domain, those fundamental

values and disvalues are constitutive of it and they ground the goodness and badness of other things in the domain through derivative value relations (again, I will have more to say about these derivative value relations later). Within a critical domain, there is always a core group of fundamental values and disvalues that generate a distinctive web of good and bad kinds of things and the evaluative statuses within these domains depend on this structure. Our own interests or goal-directed activities help explain why any particular critical domain is an important one for us to engage in. This can at least gives us a basic understanding of what it means to evaluate something "from a certain perspective" in Sosa's sense.

The connection between our interests and the relative importance of particular critical domains is never explicitly made in Sosa's work. However, seeing critical domains in this way does seem to correspond with his intentions. For example, the fact that Sosa explicitly focuses on "performances with an aim" for the sorts of evaluative structures he is imagining suggests that he does see the values within them and their importance as being explained by something like interests or goal-directed activities.[2] Part of the reason for his explicit focus on *performances* relates to his own particular virtue-theoretical conception of epistemology, but we need not follow him here (at least not yet) in order to get a grip on critical domains, generally speaking. We can have an interest in something in a way that motivates our engaging in an evaluative practice independently of whether any performances obtain. For example, we might have an interest in beautiful sunsets and thereby engage in an evaluative practice that takes such things to be fundamentally valuable. The value of beautiful sunsets might then explain

---

[2] See, e.g., Sosa (2007, pg. 23).

the value of certain meteorological or environmental conditions conducive to such

sunsets independently of whether there are any genuine performances involved.[3]

Nonetheless, the conception of critical domains outlined above seems faithful to what

Sosa would want.

One thing that *is* made explicit by Sosa, and this will be a very important point

moving forward, is that the values within critical domains are supposed to "insulated" in

some sense of the term. Consider an example he uses to express this:

> Paradoxically, one can be an adept critic within such a domain even while
> discerning in it no domain-transcendent value. Thus, someone knowledgeable
> about guns and their use for hunting, for military ends, and so on, may undergo a
> conversion that makes the use of guns abhorrent. The good shot is thus drained of
> any real value that he can discern. Nevertheless, his critical judgment within that
> domain may outstrip anyone else's, whether gun lover or not. Critical domains
> can be viewed as thus insulated... (2007, pgs. 73-4)

Here, Sosa thinks that one might make a legitimate evaluation from within some critical

domain, calling some particular thing "good" from that perspective, and yet still not think

of that entity as good or valuable at all in actuality.[4] In this way, the goodness of

something within the domain (in this case, a gunshot) is supposed to be "insulated" from

its evaluative status outside the domain. This may initially seem to be a strange way of

talking. If someone were to call a certain shot "good" and then go on to say that it is not

---

[3] This is not to say that the evaluation of certain kinds of performances will not arise further downstream in the evaluative structure. For example, someone's car-driving performance might be better or worse from that perspective if one needs to drive quickly and efficiently to a spot in order to view a sunset before dusk. The example about a possible interest in sunsets is just to point out that performances are not necessary components of certain kinds of critical domains, unlike the examples in which our constitutively-performative interests and goal-directed activities structure the fundamental values within a domain (e.g., the value of winning tennis matches).

[4] For a classic discussion of a similar point, see Foote and Montefiore (1961). Mackie (1977), who is a subjectivist about what I will call "genuine value," also talks about a more objective, standard-specific value that could easily be interpreted as the kind of domain-internal value specified here. See, e.g., pgs. 25-30.

really good, it would be very natural to think that this person is pretty straightforwardly contradicting herself.

Here, a distinction needs to be made between what I will call "domain-internal value" and "genuine value." Domain-internal value is the kind of value that an entity has within a critical domain or from a certain perspective. It is the value that accrues to something within an evaluative practice framed by certain possible interests or goal-directed activities. In Sosa's example, a gunshot exhibiting skill and accuracy would rightly be considered "good" or "valuable" from within the "world of gun marksmanship" or something of the sort. From that perspective, a shot is good insofar as it falls in line with the particular value structure that is constitutive of that domain. What is important to note is that, even if some person were repulsed by guns and any activities associated with them, she could still take up the marksmanship perspective and truly say that a skillful and accurate shot is a good one; i.e., she could truly say that it is domain-internally valuable. *Genuine* value, however, is real value.[5] Sosa's point is the following: the fact that *x* is domain-internally valuable is not sufficient for *x* to be genuinely valuable. In other words, the fact that *x* has domain-internal value does not entail that *x* has genuine value.[6]

In Sosa's example, some particular shot might be a good one from the perspective of gun marksmanship (i.e., it might be domain-internally valuable) but yet not be a good

---

[5] As to what genuine value consists in, I leave it as an open question; e.g., as to whether genuine goodness/value always needs to be qualified with respect to a particular sortal (see, e.g., Foot and Montefiore 1961), as to whether it can be reduced to goodness-for or value-for some individual or individuals (see, e.g., Thomson 2001), and as to whether it is always relative to a particular individual's desires or pro-attitudes (see, e.g., Mackie 1977). I will also leave it as an open question as to whether it applies, most basically, to individuals or to states of affairs or both.

[6] Though Sosa does not make the point, this insight goes through for disvalues as well. The fact that *x* is domain-internally disvaluable is not sufficient for *x* to be genuinely disvaluable.

or valuable thing at all. Think of the case of an assassin who has used his marksman

skills to shoot and murder some innocent victim. It may very well be that this assassin's

shot was good *qua* marksmanship but that surely does not entail that the shot was a

genuinely good or valuable thing. This thought should not be a worrisome one. Ever

since Geach (1956) has made the distinction, we have known that there are two

fundamentally different kinds of adjectival modifiers, the predicative and the attributive,

with relevantly different kinds of logical properties. Suppose we say that the marksman's

shot is *domain-internally valuable*. The point above indicates that "domain-internally"

acts attributively in the clause "domain-internally valuable." When we say, for instance,

that *x* is a big flea, we do not mean that *x* is both big and a flea. Similarly, when we say

that *x* is domain-internally valuable, we do not mean that *x* is both domain-internal (in

whatever sense) *and* is genuinely valuable.[7]

What, however, *does* "*x* is domain-internally valuable" mean? In a case where we

say that *x* is a big flea, we are saying that *x* is big *for a flea*. We are still attributing a

property to *x* that relates big-ness and flea-ness, even though that attribution does not

logically split up into two separate descriptions of an entity as being both big and a flea.

If we were to give the same explanation for "*x* is domain-internally valuable," we would

get something that cannot be correct given what has been claimed above about the case of

the assassin (where it seems as though no genuine value obtains). For example, when I

say of something that it is big for a flea, I still imply that it is a genuine member of the

kind *flea*. The very same reading of "*x* is domain-internally valuable" (were it to make

---

any sense at all) would still be implying that *x* is an instance of something *valuable*. Some cognitive dissonance would remain.

I think that we are in the best position to understand claims about domain-internal value when we understand them in the following way: When we say that *x* is domain-internally valuable, we are attributing a particular kind of *subjunctive* property to *x* that relates being domain-internal and being genuinely valuable. We are saying that *x would be genuinely valuable* for someone having certain interests or engaging in certain goal-directed activities, given that such value is not outweighed or undercut by other value considerations. In other words, we are saying that *x* would be *pro tanto* valuable for someone having those interests, where *pro tanto* values are those that might be outweighed or undercut. Take the case of a cup of coffee that is delicious and aromatic. That cup of coffee is domain-internally valuable from within the "world of coffee." Given the analysis above, that means that the cup of coffee *would be pro tanto valuable* for someone interested in or engaging in the goal-directed activity of drinking delicious, aromatic coffee. The underlying thought here is that it is a *pro tanto* good thing for one to satisfy one's own interests and achieve one's own goals, even if such achievements are not all-things-considered valuable.

For cases like the assassin, the accurate shot is a good one (i.e. it is domain-internally valuable) because it *would be pro tanto valuable* for someone having an interest in marksmanship (which, presumably, includes the assassin herself). But, of course, other values would also be relevant to the situation in a way that the *pro tanto* goodness of the shot would be undercut (or outweighed, depending on one's view). The particular accurate shot in the example is not all-things-considered valuable because the

value that would accrue by satisfying one's interests in that way is defeated by considerations of moral value, among others. The subjunctive analysis of the semantics is one way of understanding "domain-internal" as attributive where something can have a kind of domain-internal value and yet not having any genuine value. As Sosa claims, value accruing within critical domains can be "insulated" in a way that judgments about it are protected from considerations about genuine value.

To be clear, the claim that there is no direct entailment relation holding between <*x* is domain-internally valuable> and <*x* is genuinely valuable> does not mean that there are never contexts for which it would be appropriate to infer genuine value from a particular critical domain evaluation. Depending on the context of a conversation, it might just be completely obvious that a speaker *does* have an interest or goal vis-à-vis the values of the relevant evaluative domain, that there are no relevant overriding or undercutting value considerations, and that genuine value does follow. Suppose, for example, that Joe claims that he is drinking a really good cup of coffee, that it is delicious and aromatic. Given the context, it might be reasonable to infer that the cup of coffee has genuine value. Perhaps Joe has explicitly expressed his interest in delicious, aromatic coffee and, unlike the case of the assassin, it is apparent that there are no relevant defeaters. In a case like that, we should not hesitate to think that some genuine value has also been established. For Sosa's point about insulation to hold, we only need a semantic story that makes it *possible* to separate claims about domain-internal value and claims about genuine value. This is what we get with the subjunctive analysis above.[8]

---

[8] For the sake of less clunky constructions, I will henceforth omit "domain-internal" and "genuine" when the context makes the kind of value under discussion clear.

1.2: Epistemology as a Critical Domain

We are now in a position to explore some implications for epistemic evaluation in particular. Independently of Sosa's work, it is actually relatively standard for epistemologists to think of the epistemic evaluative practice as comprising something like a critical domain. At the very least, the rhetoric in the literature suggests this. Consider William Alston's way of describing epistemic evaluation: "We evaluate something epistemically… when we judge it to be more or less good or bad from the epistemic point of view, that is, for the attainment of epistemic purposes." (2005, pg. 29) Richard Foley understands ascriptions of epistemic rationality in terms of the satisfaction of a certain set of "purely epistemic goal[s]." (1993, Ch. 1) Laurence Bonjour thinks that epistemic evaluation is tied up with our distinctively "cognitive goals" or "cognitive endeavors." (1985, pg. 8) These philosophers all see epistemology and the evaluations within it as involving a certain circumscribed set of goods. We have certain interests or goals (i.e., the "cognitive" or "intellectual" ones, whatever those are) that can explain our engagement in a distinctively epistemic evaluative practice.[9]

The philosophers above are not alone in thinking this way. It seems to have permeated the epistemological enterprise. Of course, this is not to provide an argument in favor of seeing epistemology as a critical domain.[10] The argument comes when the critical domain framework is shown to be theoretically fruitful, wielding an unrivaled explanatory power. If the epistemic evaluative practice does comprise a critical domain,

---

[9] While Quine is often thought to have claimed that epistemology is a purely descriptive matter, he explicitly endorses a critical domain kind of approach to epistemic evaluation in his later work. See, e.g., Quine (1986, pgs. 664-5).

[10] In fact, there may be certain arguments that can be leveled directly against the view. I will discuss some such arguments in Chapters 3 and 4.

we should expect the following to hold: (1) that a range of interrelated fundamental and derivative values ground the distinctively epistemic evaluative statuses of the relevant kinds of things and (2) that the relative importance of the epistemic evaluative practice can be explained by appealing to an associated set of interests or goal-directed activities. In the remainder of this chapter, I will address (2) by attempting to provide an explanatory basis for the importance of the most fundamental values in the epistemological critical domain. In later chapters, I will address (1) by discussing some of the values that we should and should not expect to derive from the fundamental ones suggested here, arguing that the structure presented delivers a very intuitive and theoretically fruitful way of thinking about epistemic evaluation.

Ultimately, the claim that epistemology comprises a certain kind of critical domain is a theory about epistemology and epistemic evaluation, designed to do the sorts of things that theories normally do. Like any theory, it is invoked to explain certain phenomena, in this case, phenomena about the structure and force of epistemic evaluation. Like any theory, there will be a certain amount of moving parts within it that will make any specification more or less plausible. For instance, even in a case where two theorists both think that epistemology comprises a critical domain, if they posit different fundamental domain-internal values, they posit different theories about epistemology which might be better or worse.[11] So, one important question is the

---

[11] There is a substantive assumption lurking here about whether or not there is a single, unique evaluative practice properly known as the "epistemic" one. Alternatively, one might simply think that there are multiple, distinct evaluative practices that, through some rough association, commonly fall under the name "epistemic." If that would turn out to be the case, then there may not even be any facts of the matter that would arbitrate among what, on the surface, look like alternative theories about the epistemic domain and its most fundamental values and disvalues. Moreover, if the most fundamental epistemic values have their valuable status as a matter of conceptual truth, different theorists may simply be working with different concepts of *epistemic*, in which case the disputes involved might easily fall under what have been called metalinguistic or metaconceptual negotiations about what principles should frame our use of the concept

following: What are the fundamental values in epistemology? Many have been suggested

in the literature. Probably the most popular suggestion is that of true belief.[12]

Translating this into the critical domain terminology, many have thought that true belief

is at least one kind of rightful end to our distinctively epistemic aims. Other philosophers

have also suggested knowledge,[13] justification,[14] understanding,[15] etc., as playing this

role.

      Differences in the fundamental values posited can help individuate critical

domain theories. Another way of individuation is to count *how many* fundamental values

are posited. One recurring theme in the literature on epistemic values is the clash

between the *monists* and *pluralists* about positive fundamental epistemic value. The

monist claims that there is only one positive fundamental epistemic value and that all

other epistemic values are derivative, explained by the single fundamental one. An

example of a monist theory would be what Alvin Goldman calls "veritism," where true or

accurate belief is taken to be the sole fundamental value, purportedly explaining all of the

other epistemic values.[16] Contrarily, the pluralist claims that there are multiple

---

(see, e.g., Plunkett 2015). If that would turn out to be correct, then the framework suggested here should not necessarily be thought of as *the* one and only way to structure the epistemic evaluative practice. However, it should be thought of as a particularly interesting and useful way of thinking about epistemology, one that can do theoretical justice to many of our intuitive epistemic evaluative verdicts while also providing a substantive link between such evaluations and genuine value. It can be thought of as what Georgi Gardiner (2015) calls a "teleological approach" to the concept *epistemic*, an attempt to put the practice to good use via a particular function (also see Henderson and Horgan 2015).

[12] See, e.g., BonJour (1985, pg. 7), Foley (1993, pg. 19), Plantinga (1993a, pgs. 14-5), Goldman (2001, pg. 31), David (2001, pg. 151), Sosa (2001, pg. 51), Alston (2005, pg. 29), Lackey (2008, pg. 226), Lynch (2009, pgs. 227-8), Dogramaci (2012, pg. 514), Hazlett (2013, pg. 131), McGrath and Fantl (2013, pg. 135), Kvanvig (2014a, pg. 360), Sylvan (2014, pg. 122), among many others.

[13] See, e.g., Williamson (2000, pg. 47), Sosa (2007, pg. 88), Millar (2010, pg. 164), Hattiangadi (2010, pg. 430), and Kvanvig (2014a, pg. 360).

[14] See, e.g., DePaul (2001, pg. 180) and Kvanvig (2014a, pg. 360).

[15] See, e.g., Elgin (1996, pg. 123) and Kvanvig (2003a, pg. 200).

[16] See, e.g., his "veritistic unitarianism" developed in Goldman (2001). For other defenses of truth-monism, see Moser (1985), Latus (2000), David (2001; 2014), Douven (2008), and Ahlstrom-Vij (2012).

fundamental values within the domain, even though many other values may still be derived from those multiple fundamental ones. Sosa himself seems to advocate a pluralism about fundamental epistemic values.[17]

When formulating a critical domain theory, there are a number of considerations that need to be taken into account in determining how many fundamental values to posit. One consideration is relatively straightforward. If positing only one fundamental value, whatever it may be, cannot adequately explain all of the relevant evaluative statuses within the practice, then more will need to be posited. Pluralism would be necessary. Monists about epistemic value, however, are hopeful (and occasionally confident) that one and only one fundamental value needs to be posited to adequately explain the practice. Veritists like Goldman think that we can explain the epistemic value of any and all relevant items by appealing to the value of true belief.

Why, though, would someone like Goldman prefer monism? Why feel the need to single out one and only one value as fundamental within the domain? Why not merely locate all of the values within the domain and, say, posit them all as fundamental? In a case like that, we would not have to bother ourselves with pesky and potentially challenging derivations of other epistemic values from more fundamental ones. For example, positing multiple fundamental epistemic values instead of just one (perhaps including the value of justification) might allow Goldman to avoid Linda Zagzebski's (2003) well-known objection that his reliabilism cannot properly account for the epistemic value of knowledge over and above mere true belief.[18] If Goldman were to

---

[17] See, e.g., Sosa (2007, pg. 68) or Sosa (2009, pgs. 136-7). For other pluralist views, see DePaul (2001), Siegel (2005), Greco (2010), Gardiner (2012), and Kvanvig (2014a).

[18] The "value autonomization" solution in Goldman and Olsson (2009) actually hints in this direction. I will have more to say about this problem in Chapter 2.

expand his list of fundamental epistemic values even more to include things like states of understanding, then perhaps the sorts of worries Catherine Elgin (2004) brings out about accounting for scientific desiderata could also be avoided.

Speaking on behalf of the monist, I think the answer to these questions is that monist theories are *ceteris paribus* preferable to pluralist ones. If all else is equal in terms of explanatory power and we can deal with just one fundamental value instead of two or three, then elegance, simplicity, parsimony, etc. seem to dictate that we should. It is, of course, a substantive question for any particular monist theory as to whether it can properly account for the relevant practice.

There is, however, another consideration that might also speak in favor of monism. Monist theories might have an easier time explaining why some particular practice is an *interesting* one, one that we should find important and engaging. For example, *if* a monist theory like Goldman's properly met the explanatory desiderata, it would give us that bigger-picture story about epistemology that we have been aiming for, explaining from a single principle why epistemologists are interested in all of the various items that they are (e.g., justified beliefs, bits of knowledge, etc.). For him, such entities are epistemically valuable because they are all related, in the right ways, to the fundamental value of true belief. Insofar as true beliefs are something that we should take an interest in, Goldman's monism might be seen as *vindicating* the epistemic evaluative practice.

Consider, now, a pluralist theory positing multiple fundamental values. Unless those distinct fundamental values are related in some interesting way, it becomes a little more difficult to see why we should be motivated to care about the particular evaluative

practice in question. Suppose that I invent a critical domain and, by stipulation, it takes all and only the acquisition of true beliefs and the acquisition of money as fundamentally valuable. Although it seems like I *could* evaluate various things and states of affairs as to how well they fall in line with the values constitutive of that particular domain (given, of course, a certain amount of precisification), it is unclear as to why I should care about those evaluations. For a critical domain that purportedly has multiple fundamental values, what would help in vindicating the evaluative practice would be an explanation as to why those distinct, fundamental values are tied together in an interesting way. Unless that kind of story is given, pluralist critical domains are best avoided.

The theory I will be proposing about the epistemic evaluative practice, perhaps surprisingly, will be an attempt to stay neutral between monism and pluralism. But, I hope to do this in a helpful way. I do think that the considerations above are moving. That is, *if* a pluralist critical domain is ultimately posited, then there is a certain burden on the theorist to tell us more about what brings those multiple fundamental values together in a way that we should care about the practice. If an epistemologist says that x, y, and z are each fundamental domain-internal values, then we should expect some kind of story as to what x, y, and z have in common as interests or as the aims of certain goal-directed activities such that they are all distinctively epistemological. I will come back to this point momentarily.

The central suggestion that I want to make in this section is that the most fundamental values within the critical domain of epistemology are those structured by the constitutive goals of inquiry. Here, it will be crucial to be clear about what I mean by "inquiry," a notion typically associated with some form of asking and answering

questions. There are a number of ways to understand the association here. For my purposes, when I say that some agent S is inquiring about Q, I do not mean that S merely utters words in the form of a question Q. Someone might utter words in the form of some question Q while not really being interested in a genuine answer to Q at all, perhaps uttering Q for rhetorical purposes or for some other reason independently of getting a genuine answer. In cases like these, though there might be some sense in which S inquires about Q, S is not an inquirer in the sense I am interested in.

There are actually two different ways to understand a more genuine sort of inquiry and both will be relevant here. First, someone might be in a *state* of inquiring about something. S is in a state of inquiring about Q just in case S is consciously interested in a genuine answer to question Q (regardless of whether or not S has uttered any words in the form of Q). It might be helpful to think of such a state as initiated by a certain kind of conscious desire, a desire for a genuine answer to a question. In that sense, it is a state with what has been called a "world-to-mind" direction of fit, where the satisfaction of it would consist in the realization of its object, namely S actually obtaining a genuine answer to the question Q.[19]

Second, someone might engage in an *activity* of inquiring, which is something over and above merely being in the state of inquiring. S might consciously desire a genuine answer to Q without actually attempting to find an answer to Q or even being able to. Those who engage in the *activity* of inquiring about Q actively probe the world in order to bring about the satisfaction of the *state* of inquiry that they are in.[20] Engaging

---

[19] The idea of "direction of fit" can be traced to G. E. M. Anscombe (1957). Also see Searle (1983) and Humberstone (1992).

[20] The idea here is similar to the way that Jaakko Hintikka formulates his "interrogative model" of inquiry. See, e.g., Hintikka (1999).

in the activity of inquiring is the proper response, *ceteris paribus*, to being in the state of inquiry in the first sense. Both the state and the activity of inquiring are well-suited to describe the fundamental values of a critical domain that we should find important. *S*'s being in a state of inquiring involves having an interest in something; in this case, a genuine answer to a question. If *S* further goes about attempting to resolve her desirous state, inquiring in the second sense, then *S* engages in a goal-directed activity of bringing about the possession of a genuine answer to the question.

Each of us, to some extent or other, are inquirers in both senses of the term. We are interested in answers to certain questions and, when appropriate, we carry out the activity of finding the answers to those questions. Some of us are more inquisitive, some less. Here, the central claim is that it is the constitutive goals of inquirers, generally speaking, that compose the most fundamental values within the epistemological critical domain. To be a bit more precise, if it is possible that some agent *S* can be in a state of inquiry about question *Q* or can engage in the activity of inquiry about question *Q*, then the constitutive goals of such inquiry-activities are of fundamental epistemic value. I hope to make this view and its implications clearer below.[21]

For now, the general goals of inquiry should be a relatively plausible starting point in thinking about the values central to epistemology. Think again about an analogy with the critical domain structured by the fundamental value of delicious, aromatic coffee. Why would someone engage in an evaluative practice like that? Most naturally,

---

[21] A number of epistemologists, both implicitly and explicitly, also explain epistemic desiderata in terms of the goals of inquiry. For example, Goldman (1999, pg. 3), Lynch (2004, pgs. 15-6), and Alston (2005, pg. 31) focus in on the idea of *curiosity* in order to ground our distinctively epistemic values. The idea of curiosity that they have in mind would be most closely associated with the *state* of inquiry outlined above. Even going as far back as Carl Hempel's (1965) work on scientific desiderata, we see explicit appeals to our interests in asking and answering questions.

a person would because she, herself, is interested in what is fundamentally valuable from that perspective: instances of delicious, aromatic coffee. She is interested in participating in an evaluative practice that takes such coffee seriously. Because the values within critical domains are "insulated" from other value considerations, the evaluative practitioner within the world of coffee is able to focus in on an interest that is particularly important to her by isolating the relevant evaluative components. In this way, there is a correspondence between the fundamental values within a critical domain and what, quite naturally, would motivate someone to participate in that evaluative practice.[22]

Why do epistemologists take up the particular evaluative practice that they do? It would be natural to think that, for many, one motivation stems from an interest in the constitutive goals of inquiry and whether or not we successfully meet those goals. Epistemologists often take the goals of inquiry and the interests of inquirers seriously, as it were. Suppose, as a place-holder, that we take the goal of inquiry to be something like *information gathering*. It would be natural to think that what makes many epistemologists interested in that kind of evaluation is a drive to discover how successful we are as information gatherers. Someone like the skeptic tries to show us that we are not so successful. Alternatively, non-skeptical epistemologists formulate theories of knowledge, justification, etc. that vindicate our information gathering abilities. Skeptics and non-skeptics alike, though they deliver competing evaluations of our own situations, share a common interest in our capacities as inquirers.

---

[22] Of course, someone might engage in evaluating coffee for other reasons. Perhaps someone engages, not because she is at all interested in delicious, aromatic coffee, but because she is merely being paid to do so, giving her domain-external practical reasons to evaluate coffee in that way. Though this is certainly possible (and perhaps common), the motivation to engage in the practice would only last as long as those external, practical reasons are present.

Seeing epistemic evaluation as being structured by the constitutive goals of inquiry will be a particularly interesting way of looking at things given the discussion above about monism and pluralism. Inquiry, in the second sense, is a certain kind of goal-directed activity that is a response to being in a state of inquiry in the first sense. However, when we say that someone is in a state of inquiring, we give a general description of a state that might be satisfied in fundamentally different ways. There may be *fundamentally different kinds of inquiry* that take *fundamentally different kinds of goods* as objects to be satisfied. If that would turn out to be correct, then the theorist claiming that the goals of inquiry pick out the fundamental values in epistemology might have a pluralism about fundamental values, but she would have it in a way that could explain what ties all of those fundamental values together, grounding all of the evaluative statuses in an interesting, unified critical domain. The fundamental values, though plural, would all be goals constitutively related to inquiry. Insofar as the pluralist has this kind of explanation at hand, the advantage of monism might be neutralized.

It may seem plausible to think that there are different kinds of inquiry that would give rise to fundamentally different kinds of goods. After all, we can inquire about *whether* some fact obtains, *why* some fact obtains, *when* some event occurred, *where* something is, *which* objects have certain properties, etc. At a surface glance, we desire different kinds of things when we desire the possession of an answer for each of these various kinds of questions. If the goals of inquiry are going to act as a basis for the fundamental values of the epistemological critical domain, we need to make some sense of this. Even though inquiry, generally speaking, is a complex issue for which I could not hope to uncover all of the details here, I do wish to motivate the view that there are at

least two very important, yet potentially different, kinds of inquiry that might motivate a pluralism about fundamental epistemic values.

The first kind of inquiry that I have in mind is what I will call "propositional inquiry," where the constitutive goal of such an activity is to come to *know* some proposition *p*. The paradigm case of this kind of inquiry is a response to being in a state of inquiring about whether *p*. When we are in such a state, we have an interest in the possession of an answer to the question of whether *p* is the case. What would fully satisfy such a desire? Nothing short of knowledge. In fact, in inquiring about whether *p*, there are two conditional knowledge-goals present: 1) if *p* is true, to know that *p* and 2) if *p* is false, to know that ¬*p*. Other cases of propositional inquiry include inquiries about which objects possess some property *P* (where the goal would be to know some proposition that objects *x*, *y*, *z*, etc. possess *P*), inquiries about when some event has occurred (where the goal would be to know some proposition including an event and a time as constituents), inquiries about where something is (where the goal would be to know some proposition including an object and a place as constituents), etc. In each of these cases of propositional inquiry, the goal (at the very least) is to *know* some particularly salient proposition.

By saying that knowledge is the goal of propositional inquiry, it is important to note that it is the *constitutive* goal of that activity, not some contingent property of it. For instance, different people might engage in inquiring about whether *p* for different reasons. Alice might inquire about whether *p* merely because she is curious, Brian because it is his job, Chris because she is trying to help a friend, etc. Even though, in some sense, each of these people have different goals and aims when they engage in the

activity of inquiring about whether p, what is constitutive of each of their situations is that they have the goal of knowing whether p. Aiming at knowing whether p is what all inquirers about whether p have in common, no matter their contingent circumstances or reasons for inquiring in the first place. Henceforth, when I speak of the goals of inquiry, I mean the distinctively constitutive ones.

A natural worry about saying that knowledge is the constitutive goal of propositional inquiry would be the following: Why should we think of the goal to be a state of *knowledge*, in particular, and not, say, a mere state of true belief? If the idea is that, in inquiry, we desire to possess a genuine answer to a question, why is having a relevant true belief not enough to satisfy that desire? Having a true belief about p seems to be one way of thinking about what it would mean to "possess a genuine answer" to the question of whether p. Some philosophers do think of the goal of propositional inquiry in this way. For example, Duncan Pritchard claims that "truth is the telos of a properly conducted intellectual inquiry." (2014, pg. 114) Alvin Goldman (1999, pgs. 3-5; 2009, pg. 24), Jonathan Kvanvig (2003a, pgs. 145-50), Michael Lynch (2004, pgs. 13-5; 2009, pgs. 225-8), William Alston (2005, pgs. 4-7) and others also give a nod to the goal of mere true belief. Even Sosa suggests at one point that mere true belief is the goal of inquiry, though he wavers on this and later opts more explicitly for a knowledge-goal.[23]

I am skeptical that mere true belief is the goal of propositional inquiry for a number of reasons, but I will describe only one of those reasons here (I will discuss the truth goal in more detail in Chapter 2). Suppose that Anne finds herself interested in whether p (for simplicity, let p be some empirically knowable proposition) and that, in

---

[23] See Sosa (2001, pg. 51) and Sosa (2003, pg. 157).

the course of inquiring about whether p, she is hit on the head and comes to believe, luckily but truly, that p. Now, also suppose that this belief of hers manifests itself in the following way: She is disposed to judge that p is true, she uses p as a premise in both her practical and theoretical reasoning, she behaves as if p is true, but when she is asked why she thinks that p is true, she never makes any transitions in thought from anything that would support the truth of p outside of the fact that she just really thinks that p is true. Suppose that she even thinks that she *knows* that p is true.

Now, clearly, Anne does *not* know that p and, if she were to close her inquiry about whether p, she would do so illegitimately in some sense of the term.[24] For the sake of argument, I want to grant that this consideration, by itself, is not enough to establish that knowledge is the goal of such inquiry instead of mere true belief. As Pritchard (2014, pgs. 123-6) convincingly argues, even if an activity can only be legitimately closed by being in a certain state, that does not entail that being in that state is the goal of such an activity.[25] Nonetheless, we can still ask the question of whether Anne has indeed satisfied the goal involved in her inquiring about whether p.

Of course, given the state that Anne is in after she hits her head, she at least *thinks* that she has satisfied the goal. But, has she in fact satisfied it? I doubt it. At least part of what we want when we inquire about whether p is to adequately *track* whether p, coming to possess the correct answer in a relevantly *unlucky* way. Anne, in some sense, does possess the right answer about whether p, but she has not properly tracked whether p.

---

[24] See Kappel (2010) and Kelp (2011) for instances of the view that knowledge is an "inquiry stopper."

[25] Take, for example, the activity of producing a good cup of coffee. In order to properly close that activity, one must be in a state such that it is *established* that a good cup of coffee has been made. Does that make it the case that the goal of the activity is it *being established* that a good cup of coffee has been made? Not at all; the goal is still merely to make a good cup of coffee. What it takes to properly close an activity is not necessarily identical to the goal of that activity.

Her accurate belief is lucky in a way that detracts from her success as an inquirer. If someone were to tell Anne before she hits her head that she was about to be in the state outlined above, I think that she would be disappointed vis-à-vis her goal of inquiring about whether *p*. Thinking of knowledge as the proper goal of propositional inquiry accommodates this intuition.

Similar considerations can be given in support of the view that the constitutive goal of propositional inquiry is knowledge and not mere *justified* true belief. If we have learned anything from Gettier (1963), it is that those two things best not be conflated. It is quite possible to have a justified true belief about whether *p* while not having tracked whether *p* in a way that would satisfy the goals of inquiry. Smith, for example, might have excellent evidence for a true belief of his that some person in his office has ten coins in her pocket even though he has not *tracked* that truth. Perhaps he has only luckily stumbled upon the true belief through a combination of some misleading evidence and fortuitous environmental circumstances. If Smith were genuinely inquiring about whether someone in his office has ten coins in her pocket, I take it, again, that Smith would be disappointed with such an outcome vis-à-vis the constitutive goals of inquiry.

Whatever condition is required of knowledge over and above mere justified true belief (if knowledge even is reductively condition-governed in this way), there is a strong intuition that knowledge, not mere justified true belief, is the *ideal end* of propositional inquiry. The propositional inquirer will categorically prefer knowledge to mere justified true belief in satiating her relevant inquiry-related desires. Hence, I think that it is quite natural to think of it as the proper goal, even if falling short of that goal will sometimes

turn out to be acceptable given an inquirer's other interests.[26] The goal of the

propositional inquirer is to *find out* the truth of some proposition in whatever way she can

(whether it be via perception, memory, testimony, deductive inference, some other *a*

priori method, etc.) and, as Alan Millar (2010, pg. 164) states, "[f]inding out is nothing

less than coming to know."[27]

Thus, there seems to be a kind of inquiry-activity (i.e., propositional inquiry) such

that the constitutive goal of engaging in it is to gain knowledge of some proposition.

However, we cannot jump to the conclusion that all types of inquiry can be made to fit

this mold. There is another kind of inquiry that I will call "explanatory inquiry" where it

is far less clear that the goal is always something that can be reduced to knowledge. A

paradigmatic case of the kind of inquiry I have in mind is inquiring about *why* some fact

F obtains when it is already known that *F* obtains. Explanatory inquirers like this

constitutively seek explanations for certain phenomena. Whether such desired

explanations are of the causal-scientific variety or of the metaphysical-grounding variety

will depend upon the particular token inquiry in question. Though this is a relatively

familiar kind of inquiry, it is not immediately transparent what the exact nature of its

---

[26] If one is keen on making the distinction between a more rudimentary type of knowledge (e.g., Sosa's "animal knowledge") and a more reflective type (e.g., one requiring an agent to properly base her belief on some adequate grounds), it is the more reflective type that I take to be ideal vis-à-vis propositional inquiry. This seems to me to be exactly the right fit for the epistemological critical domain. Reflective knowledge is epistemically better than the more rudimentary cognitive states that often pass as a form of knowledge (see, e.g., Sosa 2009, pgs. 147-53). Similarly for fallibilist vs. non-fallibilist conceptions of knowledge. One might make the distinction between epistemic certainty and knowledge, arguing that the former requires some further specific feature over and above the latter. Again, if one is keen on making such distinctions, the former would be the ideal and the latter only a relative approximation of the ideal. See Chapter 2 for more discussion on relative approximations to fundamental value.

[27] Allan Hazlett (2013, pg. 3) and Christoph Kelp (forthcoming) have also suggested that knowledge is the goal of inquiry. Of course, none of this is to say that justified true beliefs (or even merely justified beliefs) will not turn out to be good or valuable with respect to the goals of propositional inquiry and hence the epistemological critical domain. It is just to say that beliefs with those properties are not the most fundamental goods. I will be arguing in Chapter 3 that there is still a special place for justified beliefs within the evaluative structure framed by the constitutive goals of inquiry.

constitutive goal would be. Suppose that you find yourself inquiring about why the sky is blue. Sure enough, what you seek is some kind of explanation. But, do you merely seek some particular bits of propositional knowledge?

Part of the reason that it is unclear whether some kind of knowledge is the goal in this context is because some genuine explanations are not wholly factive. Catherine Elgin (1996; 2004; 2009) has made this point within the context of scientific inquiry. For example, a genuine scientific explanation for why the sky is blue would appeal to certain laws that play a role in our best physical theories. We typically think that appealing to such laws and theories can provide us with genuine explanations, even though we know that the laws do not always strictly and literally hold (e.g., the ideal gas law) and we know that the theories, as a whole, are not likely to be strictly and literally true (at least in their current forms). If scientific explanations are rarely ever wholly factive, then it is not clear as to whether we can say that the goal of seeking such explanations in inquiring about why some fact *F* obtains is knowledge of some proposition or propositions. For an explanatory inquirer *S*, the most likely candidate for a knowledge-goal would be something like the following: *S* has an interest in knowing that *x* explains *F*, where *x* is some genuine explanation for why *F* obtains. Again, though, it is not clear that it can be the case that some inquirer can *know* that *x* explains *F* where the propositions appealed to in the explanans are not even strictly and literally true.

Even more, it seems doubtful that merely *knowing* that *x* explains *F* is what an explanatory inquirer is after. Take, for instance, a high school biology student who, on the basis of her teacher's testimony, *knows* that the mechanisms of natural selection explain the variable traits found in different kinds of beetles. Suppose that she also does

not have the faintest clue about *how* the explanation is supposed to work. Would the explanatory inquirer be satisfied in such a state? Seemingly not. What could be less controversially stated about the constitutive goal of explanatory inquiry is that such inquirers aim at some kind of *understanding*. In a case in which *S* is inquiring about why some fact *F* obtains, *S* has an interest in understanding why *F* in the form of grasping certain explanatory relations in a way over and above merely knowing that those explanatory relations obtain. Given what has been said above, there is then a further question of whether or not such states of understanding can be reduced to other sorts of knowledge relations.[28] For reasons to be made clear in a moment, this is a difficult question best left open here.[29]

The purpose of bringing this issue up now is that we may draw the following insight with respect to how many values should be posited for the epistemological critical domain: Given that there seem to be at least two different varieties of inquiry that someone might engage in, what the relevant goals of each type are will have bearings on the monism/pluralism debate (given an intention to use such goals in framing fundamental epistemic value). If it turns out that the goal of explanatory inquiry (i.e., something like understanding) can be fully explained in terms of propositional knowledge, then a knowledge-monism might be made more or less plausible. If,

---

[28] Elgin is not alone in thinking that understanding cannot be reduced to knowledge. Zagzebski (2001), Kvanvig (2003a, Ch. 8), and Pritchard (2010, Ch. 4) also do not think that such a reduction is possible. This directly contrasts with certain philosophers of science like Peter Achinstein (1983) and Peter Lipton (2004) who do think that the relevant sort of understanding-why within scientific inquiry can be reduced; in Lipton's case, understanding-why in science is a particular kind of knowledge about causes.

[29] Aristotle seems to have been on to a distinction similar to the one made here. For him, the senses deliver us *to eidenai*, perhaps best understood as propositional knowledge. But, *to eidenai* is not the only object of our inquiries: the senses "do not tell us the 'why' of anything – e.g., why fire is hot, they only say that it is hot." (981b11-12; cf. Hazlett 2013, pg. 22) Only our rational faculties can give us *epistêmê*, perhaps best translated as "understanding."

however, the goal of explanatory inquiry cannot be fully reduced in those terms, then a pluralism would result including, at the very least, knowledge and understanding as fundamental epistemic values. Again, though, this would be the kind of pluralism that would be able to explain why those multiple fundamental values are connected in a way that they form the basis for an interesting and unified critical domain. They would each be, in their own way, constitutive goals of inquiry.

In this project, I will take an artificially restricted interest in the goals of propositional inquiry (i.e., states of knowledge) and the sorts of goods derivatively generated from that particular kind of goal-directed activity. My reason for doing this is to be able to more efficiently focus in on the relationships among knowledge and what are typically taken to be other epistemically-value-laden properties of token beliefs (i.e., being true, being justified, etc.) and of the relationships between epistemic justification and genuine normative prescription. Though an ideal critical domain theory would specify all of the relevant details arising out of all of the relevant kinds of inquiry, considerations of space require a narrower focus. Hence, aside from positing knowledge as a positive fundamental goal in epistemology, I will remain mostly silent about the monism/pluralism debate and will, for the most part, leave the contribution of explanatory inquiry behind.[30]

---

[30] In doing so, I unfortunately also leave behind some very important epistemic desiderata. In the epistemic evaluative domain, I take explanatory inquiry's most central role to be that of grounding the value of certain properties governing the *interrelations* among the beliefs in a given set, independently of the justificatory status of those beliefs. For example, if S's beliefs are consistent with one another or if they "cohere" in the right ways, then S's beliefs might have some positive epistemic status in virtue of their relation to the goals involved in explanatory inquiry. Or, for capturing certain wide-scope Broomean desiderata about reasoning, S's making a rule-governed inference might be epistemically good in virtue of that inference's connection to understanding. A fuller treatment of these desiderata will have to wait for another time.

For now, if what has been suggested above is true, an interesting result follows. Let us call it the "value of knowledge thesis" (VKT): For every proposition p, knowing whether p has fundamental epistemic value. VKT is a natural result of the claim that the constitutive goals of inquiry explain the fundamental values within the epistemological critical domain. If it is the case that, for every possible inquiry-activity, the goal of such an activity is of fundamental epistemic value, at the very least, knowing whether p will be valuable for all semantically coherent propositions. There may be many other inquiry-related interests that would be relevant to the epistemological critical domain, perhaps some that are not reducible to a knowledge-goal; but, at the very least, the relevant interests will include knowing whether p, for all p.[31]

Before moving on, it should be mentioned that I have yet to say anything about fundamental *disvalue* in epistemology. Just as it is right to say that there are good and valuable things from the epistemic perspective, it is equally right to say that there are bad and disvaluable things as well. It is not as if, for instance, an agent who has formed a belief that p on the basis of a tarot card reading has formed her belief in a way that is evaluatively neutral from the epistemic perspective. That way of forming beliefs is genuinely *bad*. Given the goal of explaining epistemic evaluation in the way suggested, it will be important to accommodate this intuition by appealing to the constitutive goals of inquiry. This task will be most appropriate for Chapter 3 in a discussion of epistemic justification. For now, there should at least be some initial motivation to see whether the

---

[31] For readers concerned about knowability, we can still hold onto VKT even if it turns out that that some semantically coherent propositions are unknowable. A state of affairs does not become less of an ideal and does not become less valuable under the circumstances in which it is impossible to obtain. Justice for all is morally valuable even if it cannot possibly obtain. Similarly, knowing that p is epistemically valuable even if p is unknowable.

goals of inquiry can do the work needed for the epistemological critical domain.

Eventually, the theory should be judged by its fruits.

1.3: Some Initial Worries and Replies

According to VKT, for every proposition *p*, *knowing* whether *p* is fundamentally

valuable within the epistemological critical domain. An initial worry about this claim is

that there seem to be an abundance of cases of knowledge that are completely worthless

and/or trivial. Call this the worry from value-less knowledge.[32] Take, for example,

Richard, who happens to glance at a random entry in the phonebook and comes to know

that it ends in a "6." Unless Richard has some particular interest in that entry (which he

very easily might not), there is little reason to think that such knowledge would be

valuable. Is the account outlined above undesirably committed to the contrary?

It will be important to recall that, when it is claimed that something is valuable

within a critical domain, the sort of value conferred on that thing is domain-internal value

and not genuine value. The fact that something has epistemic value (in this case,

fundamental epistemic value) is not sufficient for it to have any genuine value. Thus, in

responding to the worry from value-less knowledge, it would be perfectly appropriate and

consistent to say that Richard's knowing that the relevant entry in the phonebook ends in

a "6" has fundamental epistemic value, but yet has no genuine value.[33] For all the critical

---

[32] Given the popularity of veritism about fundamental epistemic value, the most common form of the objection considered here is aimed at trivial, value-less instances of true belief. Iterations of the worry can be found in Goldman (1999; 2001), Sosa (2003), Kvanvig (2003a), Lynch (2004), Alston (2005), and Horwich (2006). An excellent summary of the positions taken with respect to it can be found in Grimm (2009). This particular issue will be discussed in more detail in Chapter 2.

[33] The application of the attributive reading of "epistemic value" here is not to deny that it is occasionally correct to take a different reading of the word "epistemic" in other contexts. Hazlett (2016) has described something like the attributive reading of "epistemic" as a "teleological use" of the word. Of course, there

domain theorist is committed to, it may even be the case that Richard's knowledge is genuinely *disvaluable*. What VKT *is* committed to is the following: Knowing whether the relevant entry ends in a "6" would be *pro tanto* valuable for a certain kind of inquirer (e.g., an inquirer about that entry), even if Richard himself is not such an inquirer.[34]

A more involved, yet related worry comes from Jason Baehr:

> Consider… much of what is peddled on the evening news (at least here in Los Angeles): the latest scoop about which Hollywood 'celebs' are sleeping with (or divorcing) each other; how these people are dressing their children; or the unsavory details of the latest sex scandal in the local schools. Here it would be an understatement to say that the knowledge in question 'lacks value'… Even from a strictly 'epistemic standpoint,' or relative to the standard of an excellent intellectual life, such knowledge is problematic: it is apparently a waste or misappropriation of cognitive capacity and resources. (2009, pg. 50)

Baehr's point here goes beyond the worry that some bits of knowledge might not turn out to have any genuine value (perhaps even being disvaluable). Because there are cases of knowledge that would be a "waste or misappropriation of cognitive capacity and resources," Baehr thinks that such bits of knowledge will not even turn out to be valuable "from the epistemic perspective."

Is this a problem for the view above? I think not. Suppose that knowledge of the proposition <Mariah Carey and Nick Cannon have recently divorced> is in question. Since knowing whether this proposition is true is the constitutive goal for a certain kind

---

may easily be other uses (e.g., Hazlett's "energetic use") that are more appropriate depending on the context of use.

[34] This claim about the *pro tanto* value of knowledge from the perspective of an inquirer is also the proper response to a case brought out by Michael Brady: "It is possible for natural curiosity about some issue to reflect a false belief or a lack of understanding, in which case we might… doubt that satisfying this curiosity has final value. For instance, suppose that I am a believer in crystal healing, and as a result I am naturally curious – I desire to know, simply for the sake of knowing – which crystals are thought to be most effective for healing which ailments." (2009, pg. 270) Surely enough, such knowledge would not have any overall value for the inquirer, but this is because satisfying such an inquiry would only be *pro tanto* valuable, a value undercut by the relevant considerations about prior false beliefs and a lack of understanding.

of inquirer (i.e., for someone who actually has an interest in Carey and Cannon's relationship status), then such knowledge has fundamental epistemic value or fundamental value "from the epistemic perspective." What, exactly, does that commit us to? Saying that this bit of knowledge has distinctively *fundamental* epistemic value is just to say that, whatever domain-internal value it has, that value is not explicable or grounded in terms of other epistemic values; i.e., it is the constitutive goal of some particular inquiry. Saying that it has distinctively *epistemic* value is just to say that it would be *pro tanto* valuable for a certain kind of inquirer. All of this is consistent with saying that such knowledge could potentially be outweighed by other value considerations, even other distinctively *epistemic* value considerations. In other words, in a similar way that the fact that such knowledge has fundamental epistemic value is no guarantee that it has genuine value, there is also no guarantee that such knowledge has an overall value, even restricting ourselves to distinctively epistemic considerations.

Why not? As Baehr suggests, agents like us have limited cognitive capacities that can be used more or less efficiently. It might be the case that my storage of certain bits of information actually impedes my ability to reach other inquiry-related goals. For instance, I might happen to be obsessed with the details of the lives of a celebrity couple to the detriment of my ability to know about other things or further understand the world in a way that would constitute a more comprehensive satisfaction of the goals of inquiry. In a case like that, the epistemic value that would accrue to my knowing that Carey and Cannon have divorced might be outweighed if I am sacrificing other inquiry-related goods that could otherwise be reached. Hence, even though knowing about their divorce would have fundamental epistemic value, the contingent circumstances of my situation

preclude it from having an overall value, even restricting ourselves to the epistemic perspective.

In a case where knowing that Carey and Cannon have divorced does not have heavy inquiry-related costs, I think we should be less inclined to say that the epistemic value of such knowledge is problematic, even if prudential or moral considerations weigh against obtaining it. Something that rings true is the following: An epistemically ideal agent S is one such that for every true proposition *p*, S knows that *p*. Among such propositions is <Mariah Carey and Nick Cannon have recently divorced.> Certainly, the kind of ideal that includes knowledge of all true propositions is one that is unobtainable for creatures like us (and is probably not even desirable), but that does not make it any less of an epistemic ideal, an ideal from within the relevant critical domain.[35] This is captured by VKT.[36]

We have been considering the following view: *x* is a fundamental epistemic value if and only if *x* is a constitutive goal of inquiry. VKT is entailed by this. Consider yet another worry about this proposal: Not all cases of knowledge arise out of an explicit inquiry in the sense that an agent, in coming to have some bit of knowledge, has acted on

---

[35] The traditional monotheistic God, who purportedly exhibits omniscience, would be an instance of an agent conforming to this particular epistemic ideal (if such a God existed, of course). However, given the view outlined above, even though this God would be ideal from the epistemic perspective, we might coherently ask a further question about whether or not omniscience is a genuinely good property to have. In saying that omniscience is an epistemic ideal, there is no presumption that it is genuinely good.

[36] One might also have a related sort of worry about the *comparative* epistemic value of certain bits of knowledge. For example, many epistemologists have the intuition that knowing that stars are formed when massive clouds of molecular hydrogen collapse due to gravity is epistemically better than, say, knowing that grain of sand *x* is so-and-so-many millimeters away from grain of sand *y* (cf. Treanor 2014). Again though, this intuition will not be in conflict with VKT. Though each bit of knowledge is purportedly epistemically valuable, it still may be the case that certain bits of knowledge are more so than others because they constitute or contribute toward a more comprehensive satisfaction of the goals of inquiry. A proponent of the view sketched above could explain why we have the intuition that the first bit of knowledge is more epistemically valuable than the second by appealing to the fact that the first contributes to a certain kind of *understanding* of the world while the second does not. For a similar case, see DePaul (2001, pgs. 173-4). See Chapters 1 and 2 in Whitcomb (2007) for more discussion on this issue.

a conscious desire for an answer to a question. For instance, in the example given above, it is not as if Richard had to be consciously interested in a certain question about the relevant entry in the phonebook in order for him to gain knowledge about it. It could have easily been the case that he merely glanced at the entry by happenstance, coming to know that it ends in a "6." Someone might worry about this in the following way: If it is supposed to be the constitutive goals of inquiry that explain why certain things are good (or bad) from the epistemic perspective, how can we can say that Richard's knowledge, which is not the result of an inquiry activity, is epistemically valuable in any sense as VKT claims?

The worry here is even more general. A very large portion of our cognitive states do not seem to be the result of an explicit inquiry activity where the constitutive goal thereof is psychologically instantiated. Our beliefs quite regularly form via subconscious cognitive processing and, yet, they each seem to be evaluable from the epistemic perspective (i.e., they either constitute knowledge or they do not, they are either justified or they are not, etc.). Given that the goals of inquiry are supposed to be doing the relevant value-theoretic explanatory work, how can we reconcile this with the view sketched above?

I think the most appropriate way to respond to this worry is to recall that the evaluations rendered within critical domains are applied *from a certain perspective*. The claim on the table is that, within the epistemological critical domain, we evaluate from the perspective of an inquirer. Given this, it may or may not be the case that what we are currently evaluating from that perspective (e.g., a cognitive state) is the product of an agent who is, in fact, such an inquirer. We can evaluate Richard's cognitive state from

the perspective of an inquirer about whether the relevant phonebook entry ends in a "6" independently of whether or not Richard himself is such an inquirer. From the perspective of someone actually interested in such information, the cognitive state that Richard finds himself in might be better or worse. If Richard *knows* whether *p*, then Richard is in the ideal cognitive state from the perspective of an inquirer about whether *p*. The fact that Richard's cognitive state is not the product of an inquiry activity is irrelevant. The same could be said about any subconsciously formed belief.[37]

What *makes it the case* that Richard's knowing whether *p* is epistemically valuable, what *grounds* its epistemic value, is the fact that knowing whether *p* is the constitutive goal of a certain kind of inquiry combined with the fact that inquiry is constitutively tied up with the epistemic evaluative practice. The facts that do the grounding work here are metaphysical facts about inquiry and about the epistemic domain, *not* facts about the actual psychological states of Richard (e.g., facts about whether he has a conscious desire for an answer to the question of whether *p*). Hence, it can be true that Richard's cognitive state is epistemically valuable even if Richard could not care less about whether *p*.[38]

---

[37] Consider an analogy with the critical domain of competitive tennis. Independently of whether or not, say, a particular tennis racket is the product of an agent who explicitly intends it to be used in the goal-directed activity associated with competitive tennis (i.e., winning tennis matches), that racket can still be better or worse from the perspective of a competitive tennis player. Imagine a tennis racket that is the product of a Davidsonian (1987) swampman-like rearrangement of molecules. That racket will still be better or worse from the perspective of a competitive tennis player even though it is not the product of any agent at all, let alone the product of an agent interested in the goals associated with competitive tennis.

[38] This explanation also serves to individuate the view outlined above from another way of reading Sosa's discussion about critical domains. Though there are many respects in which the view here is complementary to that of Allan Hazlett (2013, pgs. 256-60; 2014), he reads Sosa as a proponent of what he calls "convention-relativism" and eventually endorses the position himself. The convention-relativist grounds facts about epistemic goodness and badness in facts about the *conventions* set out by those who engage in a relevant evaluative practice (in this case, the epistemic evaluative practice). It should be relatively clear how the view above differs from this as it is *metaphysical facts about what is constitutive of inquiry* that do the grounding work relevant to any epistemic evaluative status, not facts about evaluative conventions. Whether or not Sosa is in fact a convention-relativist, I think that the view has some

As a way of leading into the discussion in Chapter 2, the final worry that I will consider here comes directly from Marian David's (2001) work on epistemic goals. David thinks that our primary positive goal from the epistemic perspective is to "obtain the truth." As he rightly notes, when someone interested in distinctively epistemic desiderata says "I want the truth," the claim involved is ambiguous: "[It] can be interpreted as saying, 'I want to believe the truth' or as saying 'I want to know the truth.'" (2001, pg. 153) Given the way in which I have associated the most fundamental epistemic desiderata with the constitutive goals of inquiry, it should be clear that I think the second interpretation is the more appropriate one. Alternatively, David, who is a committed veritist, rejects this interpretation and opts for a mere-true-belief epistemic goal.

We should consider his reasons for this. David thinks that knowledge should not be seen as a constitutive goal associated with epistemic evaluation:

> [T]he knowledge-goal is hardly ever brought up in epistemology… Isn't this a bit puzzling? I think the answer, in a nutshell is this: Invocation of the truth-goal serves primarily a theoretical need, a need that arises from the overall structure of epistemology. The knowledge-goal would not serve this need. As far as epistemology is concerned, the knowledge-goal is theoretically impotent. (pg. 153)

What need makes a knowledge-goal theoretically impotent? David goes on:

> [T]he task of epistemology pretty much reduces to the task of giving a theory of justification… However, it is usually held that, at some point, the theory of justification has to 'break out of the cycle' of epistemic concepts and provide a nonepistemic 'anchor' for justification by connecting it in some significant manner with nonepistemic concepts [like truth and belief]… It is not hard to

---

problems. Even before there were any epistemic evaluative practices around to do the sort of grounding work needed by the convention-relativist, there were still beliefs that were epistemically good or bad, epistemically justified or unjustified. Think, for example, of the very first believer. Her beliefs were either epistemically justified or they were not and they had their epistemic evaluative status independently of whether any evaluative conventions existed at the time. Also, the convention-relativist seems unpalatably committed to the view that the epistemic evaluative facts would *change* in a case where our evaluative practices change.

see how the truth-goal fits into this picture… It is generally agreed that being justified is an evaluative concept of some sort: To say that believing *p* is justified or unjustified is to evaluate believing *p,* in some sense, as a good thing or as a bad thing, as having some positive status or some negative status. The suggestion is that this type of evaluation, epistemic evaluation, is most naturally understood… as evaluating beliefs relative to the standard, or goal, of believing truth and avoiding error… Invoking the knowledge-goal would insert the concept of knowledge right into the specification of the goal, which would then no longer provide an independent anchor for understanding epistemic concepts. In particular, any attempt to understand justification relative to the knowledge-goal would invert the explanatory direction and would make the whole approach circular and entirely unilluminating. After all, knowledge was supposed to be explained in terms of justification and not the other way around. (pg. 153-4)

Granting for the sake of argument the overstated claim about the task of epistemology, I think we can reconstruct David's argument as an interesting challenge to the view sketched above. It would go something like this:

1) Justification, although crucial to the epistemic evaluative practice, is most naturally understood as a derivative epistemic value, explained by the more fundamental goal of "obtaining the truth" (which can be interpreted as either obtaining true beliefs or obtaining knowledge).
2) If we explain the epistemic value of justification in terms of a knowledge-goal, our epistemological theory would turn out to be circular and unilluminating.
3) We do not want our epistemological theory to be circular and unilluminating.
4) Hence, we should not posit a knowledge-goal in epistemology, at least with respect to explaining the value of justification; a true-belief-goal would be more appropriate.

Thus, David opts for a veritistic approach to epistemic value.

Let us grant premise one for now (I will discuss it in more detail in the coming chapters). Premise three is clearly true and the argument does seem to be valid. That leaves us with premise two. There are a number of assumptions packed into this premise that are worth rooting out. Why does David think that a knowledge-goal would make our epistemological theory circular and unilluminating? The idea seems to be that, *if* we give an explanation about epistemic justification in terms of some knowledge-related desiderata, the traditional epistemologist who analyzes knowledge on the model of

justified true belief (plus some anti-Gettierizing component) would thereby be explaining justification in terms of knowledge and, in turn, explaining knowledge in terms of justification; hence, circularity.

There are two important points to bring up about David's worry here. First, he conflates two seemingly independent tasks: (1) the task of giving a metaphysical analysis of property *x* in terms of property *y* and (2) the task of explaining the value of property *x* in terms of the value of property *y*. For a derivatively valuable property within a critical domain, performing the second task (e.g., explaining the value of justification in terms of the value of knowledge) does not thereby commit one to performing the first. That is, explaining the value of justification in terms of the value of knowledge does not commit one to giving a metaphysical analysis of justification in terms of knowledge.

Take an example from the critical domain of competitive tennis: Within that domain, the property of one's tennis racket being strung at a tension between 50 and 60 pounds is derivatively valuable. Why is it valuable? It is valuable because it is the ideal tension range conducive to winning tennis matches which is the most fundamental value in that domain. The range allows for enough elasticity to utilize a wide variety of spins when striking the ball and yet also allows for enough tension to generate speed in order to overpower one's opponent when appropriate. Now, because we have explained the value of that tension range in terms of the fundamental value of winning tennis matches, are we thereby committed to giving a metaphysical analysis of that property in terms of winning tennis matches? Of course not! *What it is* for a racket to be strung at a tension between 50 and 60 pounds can be specified independently of the ultimate goal of winning tennis matches; presumably, we would need to explain the property's necessary connection to

certain physical facts about tensile strength, units of pressure, net forces, etc., but no appeal to winning tennis matches would be needed.

Something similar could be said about the epistemic value of justification in relation to the value of knowledge. Even if someone were inclined to explain the value of being justified in terms of the value of knowing, there is no necessary commitment to metaphysically analyzing the property of justification in terms of knowledge. Again, these are two independent tasks. This consideration undermines the sort of circularity posited in premise two. A theorist who explains the epistemic value of justification in terms of a knowledge-goal would be free to give an analysis of knowledge in the traditional style of JTB+, so long as a further metaphysical analysis of justification does not appeal back to the concept of knowledge. If this were done, we would have a genuinely illuminating theory about epistemic value with no vicious circularity.

The second point is perhaps just as important. Even if explaining the value of justification in terms of knowledge *did* commit one to metaphysically analyzing justification in terms of knowledge, vicious circularity need not result. Someone attracted to a Williamsonian-style, knowledge-first approach might take the concept of knowledge to be the most basic unit of analysis in epistemology, owning up to the possibility that there simply is no genuine reductive account of knowledge in terms of JTB+ to be given.[39] Jonathan Jenkins Ichikawa (2014) takes exactly this kind of approach. If his project were successful and he were to further explain the value of justification in terms of the value of knowledge, the fact that he takes knowledge as the basic unit of metaphysical analysis generates the result that his account would not be

---

[39] See, in particular, Williamson (2000, pgs. 27-48).

viciously circular in the way that David imagines . Hence, the second premise in the

argument above is weak on multiple counts. It is appropriate to conclude that David does

not present a genuine challenge to the view in which knowledge is a constitutive goal

associated with epistemic evaluation.

Chapter 2: The Value Problem of Knowledge and the Value of True Belief

In this chapter, I formulate a version of the so-called "value problem of knowledge," using it as a springboard for further discussion about the epistemological critical domain. Like a number of others, I find myself dissatisfied with some of the more traditional formulations of the problem and think that some added clarity is needed in order to draw any fruitful conclusions.[40] Ultimately, I use what I think is the most legitimate, modal form of the problem (along with my own solution) as an abductive argument in favor of the inquiry-centered critical domain proposed in Chapter 1. In the first section of this chapter, I discuss a number of ways that one might conceive of the value problem of knowledge, eventually settling on a formulation that utilizes the notion of epistemic value. In the second, I discuss a variety of what I call a "bottom-up approaches" to this problem. These approaches are attempts to explain the added epistemic value of knowledge in terms of the more fundamental epistemic value of true belief. Although the bottom-up approach has been extremely popular in the literature on epistemic value, I present formidable difficulties (not all of which are novel) for even the most promising extant bottom-up solutions and, in the final section, I argue that my own approach meets the explanatory desiderata in a more satisfying way.

2.1: Searching for a Value Problem

Many epistemologists have had the intuition that acquiring knowledge about some topic is more valuable than merely coming to a true belief about it. Knowledge just

---

[40] See, in particular, Baehr (2009) and Ridge (2013).

seems better than mere true belief. Historically, the intuition can be traced as far back as Plato where, in the *Meno*, Socrates contemplates ways of justifying it. Why, he asks, would it be better to *know* the way to Larissa than to merely have a *correct judgment* about it? (97a-97d) Either way, someone desiring to get to Larissa would seemingly have the cognitive means of getting there. Plato's dialogue constitutes the first formulation of the value problem. There is supposed to be a value *problem* of knowledge because it is not exactly clear how to explain *why* knowledge is supposed to be better than mere true belief, especially given that the practical benefits of their attainment seem equal. We can formulate an initial gloss on the value problem of knowledge: (VP) Why is knowledge more valuable than mere true belief?

Socrates attempts to construct his own solution by drawing an analogy to the statues of Daedalus. (97d-98b) Just as the statues are "tethered" in a way that they are not easily lost or stolen, knowledge is cognitively tied down such that it will not easily "run away from a man's mind." According to Socrates, even though knowledge and mere true belief have an equal practical value (insofar as they are both retained through time), knowledge is still better than mere true belief because it is more *cross-temporally stable*. Imagine that you are trying to get to Larissa and that you merely correctly believe that the road in front of you will take you there. Socrates's point is that you might easily find yourself in a situation that would cause you to give up that correct belief. Perhaps, while on the way to Larissa, the road takes a sharp turn in a way that would cause you to doubt that you are taking the correct path. According to Socrates, you would not so easily abandon your belief if you instead *knew* that the road in front of you would take

you to Larissa. Knowledge has more value than mere true belief because it is more likely to stick around in the face of adversity.

It is probably not so surprising that few philosophers have found Socrates's explanation compelling, especially given the resurgence of the value problem in the literature in the past twenty or so years. Among other considerations, Socrates seems to ignore how easily and how often human knowledge is lost. We are forgetful creatures and we lose a good amount of our knowledge through time. The human cognitive system, like any imbedded in a biological organism, is subject to degradation and failure. More, we are often *creative* in our cognitive lives in a way that disrupts our knowledge base. We subconsciously change our memories and selectively retain the ones that support the propositions that we want to be true, even if we once knew otherwise. Many cases of knowledge are not "tethered" in a way that they would be immune to such loss, at least no more so than certain cases of mere true belief that do not constitute knowledge. When ordinary knowledge is counterfactually compared to, say, *dogmatic* true belief, it even seems *less* cross-temporally stable than mere true belief because the merely true beliefs in question are so deeply ingrained that almost nothing could unseat them.

Nonetheless, for many, the original intuition persists. Knowledge still seems better than mere true belief. Socrates's approach will at least need some beefing up if it is to be satisfying. One attempt to advance the suggestion about cross-temporal stability has been made by Timothy Williamson: "Present knowledge is less vulnerable than mere present true belief to *rational* undermining by future evidence… If your cognitive faculties are in good order, the probability of your believing *p* tomorrow is greater

conditional on your knowing *p* today than on your merely believing *p* truly today." (2000, pg. 79) Like Socrates, Williamson thinks that there are cross-temporal differences between knowledge and mere true belief. Call the following "the Williamson conditional:" If *S*'s cognitive faculties are in good order, then *S* is more likely to retain a true belief that *p* through time in cases where *S* knows that *p* than in cases where *S* merely truly believes that *p*.

What is supposed to make the Williamson conditional an improvement on Socrates's solution is the appeal to the language of probability (as opposed to the mere metaphor of "tethering") and Williamson's making the cross-temporal stability of knowledge conditional on one's cognitive faculties being in "good order." Williamson recognizes that knowledge will not be more cross-temporally stable as counterfactually compared to cases of dogmatic true belief. However, he thinks that this does not undermine a claim about *rational* cross-temporal stability as cases of dogmatic belief are instances of cognitive faculties that are *not* in "good order."[41] Williamson would agree that, if *S* has a dogmatic true belief that *p*, *S*'s belief will not likely be lost even in the face of conflicting (yet misleading) evidence. Yet, if *S* were *rational*, such misleading evidence *would* undermine *S*'s true belief that *p* in a way that makes it more likely to be lost than if *S knew* that *p*. If *S* knows that *p* as opposed to merely truly believing that *p*, the body of possible evidence that could rationally undermine *S*'s belief that *p* would be smaller. For example, if Socrates *knows* that the way to Larissa is via Elm Street, then

---

[41] Jonathan Kvanvig (2003a, pgs. 14-5) presses Williamson on this point, arguing that Williamson does not pay sufficient enough attention to the pragmatic dimension in belief fixation. According to Kvanvig, certain dogmatic beliefs have a survival value in a way that they count as products of a cognitive system that is in "good order." Though I am sympathetic to Kvanvig's point (I will have more to say about the pragmatic dimension of belief fixation later in this chapter and also in Chapter 4), I will not focus on it here.

presumably Socrates has good reason to think that the way to Larissa is via Elm Street. It would then take a stronger set of misleading evidential considerations for Socrates's belief to be rationally undermined than it otherwise would be if he merely truly believed that the way to Larissa is via Elm Street.

The biggest worry about Williamson's conditional is that, even if it is true, it does not seem to fully capture the explanatory desiderata. For instance, since Williamson makes the cross-temporal advantage of knowledge over mere true belief *conditional* on one's cognitive faculties being in "good order," it still leaves open the possibility for cases where knowledge has no more cross-temporal stability than that of mere true belief (i.e., those in which the agent's cognitive faculties are not in "good order"). Then, like Socrates's analogy with "tethering," the conditional cannot, by itself, explain why someone would still think that knowledge is better than mere true belief in those sorts of cases. And many still do think that knowledge is better.

Take, for example, a particular case of a dogmatic belief that is true by luck. It would seemingly be better, in at least some sense of "better," to have knowledge. In a case like that, even though Williamson's conditional would come out true, the dogmatic true belief may still be just as cross-temporally stable as any knowledge counterpart and we might still rightfully wonder why knowledge is supposed to be better. Perhaps Williamson would say that, even though knowledge is not more cross-temporally stable simpliciter than the dogmatic true belief, it is still more *rationally* cross-temporally stable. However, given the practical thrust of cross-temporal explanations, we are left wondering why we should care more about *rational* cross-temporal stability than just cross-temporal stability simpliciter. It may be a perfectly sensible thing for someone to

point to the rationality of knowledge-constituting beliefs in order to explain their
superiority to mere true beliefs, but then practical, cross-temporal stability considerations
will not be doing any of the real explanatory work.

Some philosophers have taken the failure of practical explanations like
Williamson's to indicate something important about the real desiderata behind the value
problem of knowledge. Given the above, the best that any practical explanation can do is
explain why knowledge is *contingently* better than mere true belief. Suppose that *S*
merely truly believes that *p* and that we are considering the following proposition: <*S*'s
knowing that *p* would be better than her merely truly believing that *p*.> This proposition
may very well be true given *S*'s actual circumstances, whatever they are, even if we
restrict ourselves to purely practical benefits. The nearest possible worlds in which *S*
alternatively comes to know that *p* might be ones in which the practical utility of her
relevant cognitive state is increased. Yet, even so, the proposition would only be
contingently true. If *S*'s circumstances were to fit the right mold, the alternative
knowledge-worlds might be ones in which the practical utility of her mere true belief
would be the same as (or even greater than) the knowledge counterparts. Hence, for
many theorists about the value problem, focusing exclusively on practical utility would
deliver an unacceptable result. Rightly or wrongly, it would imply that there are worlds
in which an agent's knowledge is not better than mere true belief.

Take, for instance, Jonathan Kvanvig's canonical way of seeing value problem:

> [A] satisfactory answer to the question of the value of knowledge will need to
> explain why knowledge is, by its very nature, more valuable than its parts. It will
> not be enough, for example, to show that sometimes or in some places knowledge
> is more valuable than its parts. Instead, we will need to show that no matter what
> the world happens to be like, knowledge is more valuable. (2003a, pg. xii)

If you will forgive Kvanvig's liberal use of the word "parts" (i.e., knowledge is not literally a mereological sum of true belief and whatever else), the idea here is that a central desideratum in posing the value problem is to provide an explanation for why instances of knowledge are *non-contingently* more valuable than instances of the traditional conditions on knowledge, including that of mere true belief.[42] Since practical explanations only seem to offer contingent explanations, they are rejected by Kvanvig as being fully acceptable.[43]

For a precise formulation of the present intuition, however, we will need to be more careful than Kvanvig is as he seems to conflate two distinct desiderata. In the quote above, he first requires an explanation for why knowledge is "by its very nature" more valuable than mere true belief. This requirement, if legitimate, would be one on a specification of the *essence* of knowledge. If it were imposed, then any adequate real definition of *what knowledge is* would thereby need to specify the features of knowledge for which one could explain why it is more valuable than mere true belief. An analysis of knowledge's nature that did not give us such features should be rejected. Many philosophers have taken a requirement like this one to heart. For example, such a requirement is a motivating force in a number of virtue theoretic accounts of knowledge including those of Linda Zagzebski (1996; 2000; 2003), Wayne Riggs (2002), and John

---

[42] In formulating the value problem, since Kvanvig requires an explanation for why knowledge has a greater value than its "parts," he does not merely want an explanation for why knowledge is better than mere true belief. He also wants an explanation for why knowledge is more valuable than anything that falls short of knowledge, including cases of, say, justified true belief. This alternative way of posing the problem corresponds with what Duncan Pritchard has called the "secondary value problem." (2010, pgs. 5-7) That is, even if one has solved the "primary value problem" (i.e., explaining why knowledge is better than mere true belief), that does not thereby guarantee that one has explained why knowledge is more valuable than *anything* that falls short of it. For the purposes here, I will be focusing on the comparative value of knowledge and mere true belief and will only occasionally discuss further desiderata. I will discuss the value of justification in detail in the coming chapters.

[43] See Kvanvig (2003a, Ch. 1).

Greco (2003; 2010). For them, if knowledge, by its very nature, is some kind of *virtuous* true belief, we might have a viable explanation of its non-contingent, superior value to mere true belief. Flipping an essence requirement on its head, Williamson (2000) denies that knowledge is reductively analyzable. As he would say, the importance of knowledge "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*?"[44] (pg. 31)

However, when Kvanvig further specifies exactly what he has in mind for his requirement, he slips to a less restrictive, modal formulation of the value problem. He moves to a formulation in which we merely need to explain why knowledge is better than mere true belief "no matter what the world happens to be like." In other words, we merely need to explain why knowledge is *necessarily* better than mere true belief. Requirements about *essence* and requirements about *necessity* are not one and the same. There are many necessary features of things, even metaphysically necessary features of things that would not be included in a specification of the essence of those things. For example, it is necessary of Socrates that he is a member of the singleton set containing only Socrates, yet we would be strained to say that being a member of that set is part of the essence of Socrates, part of *what it is* to be Socrates (see Fine 1994). Hence, it seems to be an open possibility that some necessary feature of knowledge could explain why it is non-contingently more valuable than mere true belief while it is also true that such a feature would not be included in a specification of the essence of knowledge. In that case, a modal requirement would be met while an essence requirement would not be.

---

[44] For a criticism of this line of reasoning about analyses and knowledge, see DePaul (2009).

For the purposes here, I will begin with the less restrictive, modal requirement on the value problem instead of an essence requirement. Given some of the ways that Kvanvig subsequently talks about the value problem, there is good reason to think that he really meant to endorse the less restrictive, modal requirement all along and that his talk of knowledge's "nature" was an unfortunate slip.[45] When I discuss my own solution to what I will take to be the most relevant form of the value problem (in section 3 of this chapter), I will have more to say about the issues of essence and necessity. Consider, then, another gloss on the value problem that would do justice to modal requirements like Kvanvig's: (VP1) Why is it the case that, for any proposition p, tokens of knowing that p are *necessarily* more valuable than tokens of mere true belief that p? This is a much more precise way of thinking about the value problem than VP and it puts further constraints on what a good solution would look like, ruling out contingent, practical explanations. Built into VP1 is the modal intuition that knowledge is *necessarily* better than mere true belief along with a categorical constraint that it is *everywhere* and *always* better than mere true belief, no matter the proposition in question.

However, as things stand, someone might easily worry that VP1 is too strong. Jason Baehr (2009) gives us an apparent counterexample:

> [A]s a matter of fact, knowledge is *not* always more valuable than [mere] true belief. This is evident… in cases in which a demand for knowledge rather than mere true belief would require the forfeiture of certain other important goods. For instance, if trying to flee a certain foreign location in the face of some impending catastrophe, I might do better simply to trust my hunch (which happens to be accurate) about the appropriate way out than to stick around and do what it takes to acquire knowledge about this matter (e.g. find a map, get directions from a reliable passerby, etc.). This suggests that knowledge is sometimes *less* valuable than mere true belief. (pg. 49)

---

[45] See, e.g., some of the clarifications he gives on pg. 5.

Whether or not the particular case that Baehr is imagining is a genuine case of mere true belief (one might wonder whether "trusting a hunch" in the case of an emergency is a genuine form of believing), the real point of his purported counterexample is to show us that particular tokens of knowledge often have *costs* associated with their attainment; those costs might very easily outweigh whatever value a certain bit of knowledge brings to the table over and above mere true belief. Knowledge is often difficult to obtain as compared to mere true belief and, in its attainment, there may be negative effects that it would be better, all things considered, to avoid. Given the contingent circumstances of a situation, we may be better off simply having a lucky true belief instead of knowledge.

There is a move available to the proponent of VP1 and Baehr himself is aware of this. Someone who really thinks that knowledge is better than mere true belief in some necessary way will not think that the purported counterexamples are genuine. A proponent of the modal intuition might claim that the only reason Baehr is able to imagine circumstances in which mere true belief is more valuable than knowledge is that he is not considering legitimate counterfactuals in which knowledge is substituted for mere true belief while all other value-laden factors remain constant. He is imagining competing scenarios in which extra value-laden facts are stacked against the value of obtaining knowledge. That is, Baehr's proposed scenarios are not cases in which mere true belief is better than knowledge, *ceteris paribus*. Responding to Baehr would then suggest the following precisification in the value problem: (VP2) Why is it the case that, for any proposition *p*, tokens of knowing that *p* are *necessarily* more valuable than tokens of mere true belief that *p*, *ceteris paribus*?

By adding a *ceteris paribus* clause, the modal intuition is safe from Baehr's imagined scenarios. For a proponent of VP2, what makes the purported counterexample illegitimate is that there are value-laden facts that are built right into the counterfactual scenarios being imagined, facts about the situation that are independent of the relevant tokens of knowledge and mere true belief. Baehr is imagining a case in which *time* is a relevant value-laden factor and, since knowledge typically (but not necessarily) takes a significant amount of time to be acquired, it seems less optimal than luckily trusting an accurate hunch in a dangerous situation. A more legitimate case to examine would be the very same scenario but where the time it would take to come to the relevant cognitive states would be equal. If we were again to come to the verdict that it would be better to have mere true belief instead of knowledge, then we would have a more genuine counterexample. But, of course, this is exactly what a proponent of the intuition behind VP2 would deny.

At this point, it is important to note that the sort of value attributed to knowledge over and above mere true belief in VP2 is some kind of *genuine value*. We are invited to think that it is necessarily better, in some genuine sense of "better," to have knowledge instead of mere true belief (*ceteris paribus*). While I will end up being sympathetic to a certain other modal specification of the value problem, I think that the intuition behind VP2 is problematic for reasons that have not been fully appreciated. One implication of the modal intuition about genuine value is that knowledge is, everywhere and always, genuinely valuable in at least some sense of the term (see footnote 5 in Chapter 1). Something cannot be more valuable than something else unless that thing itself has value. If that is true, then there is a further assumption involved in VP2 that tokens of knowing

that *p* necessarily have genuine value. A proponent of this intuition might further believe

that knowledge has genuine value *intrinsically*, being valuable in virtue of the very kind

of thing that it is.

These further assumptions associated with VP2 are problematic for reasons

briefly discussed in the final section of Chapter 1. There seem to be authentic cases of

worthless or *valueless* knowledge. If there are cases of worthless knowledge, then there

are cases of knowledge that are not genuinely valuable. And if there are cases of

knowledge that are not genuinely valuable, then, of course, knowledge cannot be

necessarily valuable (and cannot be intrinsically valuable). VP2 would be thus

undermined. Suppose, for example, that you were offered some bit of knowledge about

how many grains of sand there are in some arbitrarily circumscribed section of a random,

deserted beach. Would you take it? Would it be genuinely valuable, in any sense, to

acquire such knowledge? I am inclined to say "no." Such knowledge is worthless.

Perhaps the defender of the intrinsic value of knowledge would respond in the

following way: You are only inclined to say "no" to this bit of knowledge because you

are imagining a scenario in which the value of knowledge is outweighed by other value-

laden factors. For instance, the knowledge you would be obtaining would be taking up

valuable cognitive space that could otherwise be used for something more important.

That bit of knowledge still has value, it is just outweighed by value considerations

independent of the knowledge itself.[46]

---

[46] The kind of reply imagined here is in line with the way Kvanvig (2003a, Ch. 2), Lynch (2004, pgs. 15-9), Horwich (2006), and Brady (2009) defend the intrinsic value of true belief. Since there seem to also be genuine cases of worthless true belief, the considerations that I outline below would apply, *mutatis mutandis*, to that position as well. One extra complication arising about the view that true belief is intrinsically valuable is that truth is occasionally seen to be the "aim of belief" or the "aim of cognition" in some important sense (see, e.g., Williams 1973, Humberstone 1992, Railton 1994, Velleman 2000, Wedgwood 2002 and 2013a, Shah 2003, Shah and Velleman 2005, Alston 2005, Steglich-Petersen 2006

I remain unconvinced. We can imagine the scenario above to be one in which we are simultaneously offered a sort of "cognitive enhancement" whereby we are assured that the extra bit of knowledge that we will be obtaining will not disrupt any of our other cognitive states and will not be taking up any extra resources that could be used otherwise. We can even imagine the scenario to be one in which it takes no time to gain the relevant knowledge and to be one in which we are assured that such knowledge will not waste our time in the future by unwillingly forcing its way into our conscious awareness. If knowledge truly had intrinsic value, we should be inclined to say "yes" to such a scenario where there are no extrinsic drawbacks. Yet, I am indifferent to such knowledge and I do not think that I am being irrational for being indifferent. Again, it seems worthless.

At this point, I think that it is a fair dialectical move to place the argumentative burden on the side of the proponent of knowledge's intrinsic value. If I were to gain the sort of knowledge offered in the above scenario, it would not affect my life in any way, nor would it affect anyone else's. Why think that such knowledge is genuinely valuable *at all*? What reasons are there for thinking that someone like me who is indifferent to such knowledge is making a mistake? Until such reasons are given, the assumption that knowledge always has genuine value is highly questionable. This undermines the formulation of the value problem in VP2 which assumes that knowledge, everywhere and always, has such value.

---

and 2009, Whiting 2010 and 2013, McHugh 2011 and 2012, and Engel 2013). I am skeptical, however, that there is a sensible way of interpreting the claim that belief "aims at truth" that would support the view that true belief is intrinsically valuable. See Grimm (2008) and Hazlett (2013, Chs. 6-8) for more on this issue. I will have more to say about the "aim of belief" literature in Chapter 4.

What I think that these considerations tell us about the value problem is that a legitimate categorical and modal formulation of it cannot be found by appealing to knowledge's necessary, *genuine* value (in whatever sense) over and above mere true belief. Sometimes knowledge is just not genuinely valuable and, hence, is not more genuinely valuable than mere true belief. Mandating an answer to VP2 is a call for an explanation of a phantom phenomenon, akin to demanding an answer to the question "Why is the earth flat?" If the presupposition in a call for an explanation is false, then there is no real need to answer the call for an explanation. In fact, there is no possible way *to* answer the call for an explanation.

One way of responding to such considerations would be to give up on a modal formulation of the value problem altogether, perhaps contenting oneself with contingent explanations like Williamson's. Although I am friendly to the idea that practical considerations can explain the differences in *genuine* value between knowledge and mere true belief in the relevant cases, I do not think that we should give up on a general, modal formulation quite yet. There is still an interesting formulation to grapple with, one that might explain why many philosophers have found a more general, modal intuition to be so compelling. A number of philosophers who engage with the value problem formulate it, not in terms of genuine value, but in terms of "epistemic value."[47] Consider, then, an epistemic formulation of the value problem of knowledge: (EVP) Why is it the case that, for any proposition *p*, tokens of knowing that *p* are necessarily more *epistemically* valuable than tokens of mere true belief that *p, ceteris paribus*?

---

[47] See, among others, DePaul (2001), Pritchard (2010), Chrisman (2012), and Sylvan (2014). Unfortunately, DePaul and Pritchard in particular fail to distinguish between genuine value and domain-internal value in the way that I have here. This complicates things with respect to how well their own views can be mapped onto the framework that has been set up.

It is worth noting that a *predicative* reading of "epistemically valuable" in EVP will render it just as problematic as VP2. That is, if we interpret the claim that *x* is epistemically valuable as meaning that *x* is both of the epistemic variety (in whatever sense) *and* is genuinely valuable, then EVP will have built in the very same faulty assumption that knowledge is, everywhere and always, genuinely valuable.[48] Given the considerations above, we should avoid such a reading. A more promising way of interpreting EVP is to take an *attributive* reading of "epistemically valuable" in the manner discussed in Chapter 1 (see pgs. 11-5). On such a reading, EVP would be assuming that knowledge is, everywhere and always, more epistemically valuable than mere true belief but it would *not* do so in a way that assumes knowledge to be, everywhere and always, genuinely valuable. Hence, worries about worthless knowledge would not apply.

Is knowledge necessarily more epistemically valuable than mere true belief in the attributive sense? I think that nearly all epistemologists would say that it is. Compare two possible relations that some agent *S* might have to a true proposition *p*: (1) *S* might believe that *p* and (2) *S* might know that *p*. Which would be better from the epistemic perspective? I take it as relatively uncontroversial that, *ceteris paribus*, knowing would be better. Further, knowledge would seem to be *necessarily* epistemically better than mere true belief, regardless of whether or not any genuine value accrues to either. An attributive reading of EVP is a call for an explanation of that fact. Is it somehow explained by the nature of knowledge? Is it explained by the nature of epistemic

---

[48] Kvanvig (2014a) seems to have explicitly endorsed a predicative reading of such claims, aiming his inquiries at "values or goods that are epistemic in character." (pg. 352) Pritchard (2010) also seems to have taken a predicative reading of claims about epistemic value, though he is less explicit about this and only later comes to endorse an attributive reading (see, e.g., Pritchard 2014).

evaluation (i.e., the epistemological critical domain)? Is it explained by something else?

For the remainder of the chapter, I will take EVP to be the most legitimate general and

modal formulation of the value problem. It is a question that seems deserving of an

answer. If it is not a legitimate question, then we are at least owed some explanation as

to why it is not. Requesting an answer to EVP by no means entails that an answer must

be extraordinarily deep or labyrinthine (in fact, I take my own answer to be quite simple)

but any theorist about epistemology should have *something* to say about it.

2.2: Bottom-Up Approaches to the Value Problem

In this section, I discuss a number of possible solutions to the value problem of

knowledge, each of which are instances of what I call a "bottom-up approach." The

bottom-up approach takes tokens of *true belief* to be fundamentally valuable from the

epistemic perspective and then builds the value of knowledge up "from its parts," as it

were. Less metaphorically, proponents of the bottom-up approach begin theorizing about

epistemic value with the fundamental value of true belief and then attempt to explain

knowledge's added epistemic value through the value of its other essential properties

(e.g., being a certain kind of justified belief, being an "achievement," etc.) which

purportedly have a certain kind of derivative epistemic value (in a sense to specified).[49]

For the bottom-upper, true belief is fundamentally valuable and knowledge is essentially

TB+ (i.e., a certain kind of true belief). Purportedly, the extra properties needed in order

to turn a mere true belief into knowledge will allow us to solve the value problem.

---

[49] There are some who might worry about the bottom-up approach because it implies that belief is a
necessary condition on knowledge. I, myself, do not find the purported counterexamples to belief's
necessity (see, e.g., Radford 1966) particularly convincing and I will not discuss this worry here. I only
mention it as another possible objection.

For considerations of space, there is no way to discuss all of the possible constructions of a bottom-up solution to the value problem. However, I do discuss three distinct bottom-up approaches that will plausibly allow us to draw some general conclusions. I discuss, at length, a process reliabilist solution from Alvin Goldman and Erik Olsson, Kurt Sylvan's non-instrumentalist veritist solution, and John Greco's virtue reliabilist solution. I argue that each of these purported solutions faces formidable difficulties in a way that should cause us to doubt whether there is a satisfactory bottom-up solution available to us. Unfortunately, a good amount of the literature on the value problem is systematically ambiguous. Many commentators simply do not distinguish between VP1, VP2, predicative readings of EVP, and attributive readings of EVP.[50] On the surface, this is somewhat understandable. It is easy to work one's way into the problem through a very vague intuition that knowledge is better than mere true belief. However, as we have seen, some interpretations of the problem are more legitimate than others and, since I take an attributive reading of EVP to be the most uncontroversial and pressing form, I frame all of the possible solutions here in terms of whether or not it can meet that particular challenge.

Goldman and Olsson's (2009) approach to the value problem is a paradigmatic starting point for what I am calling a "bottom-up" solution. After all, it was Goldman who coined the term "veritism" in order to designate a cluster of value-theoretic views, including his own, claiming that true belief is the only fundamental epistemic value while all other epistemic values are derivative.[51] A proponent of veritism is forced into

---

[50] Kurt Sylvan is a notable exception here and he makes it clear that he has something like an attributive reading of EVP in mind.

[51] For more on veritism, see Goldman (1999; 2001). Veritism is what I have called a "monist" theory in Chapter 1 (see pgs. 18-21).

attempting a bottom-up solution to the value problem of knowledge because of her self-imposed restrictions on fundamental epistemic value. If true belief is the only fundamental epistemic value and knowledge is supposed to be epistemically better than mere true belief, then there is a need to explain that added value of knowledge in terms of its other properties which are claimed to be derivatively valuable relative to true belief.

It will be useful, at this point, to be a bit clearer about the way I am thinking of derivative value relations within critical domains. In the previous chapter, I described fundamental values as *grounding* the derivative ones in a way that they metaphysically *explain* their status as valuable things. For the purposes here, I will specify a way of connecting derivative values to more fundamental ones in terms of the potentially less strained notion of ontological dependence.[52] I take the following dependence condition to hold on values within critical domains: If *x* and *y* are valuable with respect to domain *D* and the domain-internal value of *x* ontologically depends on the domain-internal value of *y* in an asymmetrical way, then *x* is derivatively valuable with respect to domain *D*. A domain-internal value will be derivatively valuable when its status as a value asymmetrically depends on the domain-internal value of some other kind of thing. We can return to Sosa's "world of coffee" for an example of this kind of dependence relation: Since the coffee-related value of a certain coffee machine ontologically depends on the

---

[52] One reason that the notion of dependence may be less strained than the notion of grounding is that it does not invoke the potentially questionable notion of metaphysical explanation. In the way that some metaphysicians think of grounding, the kind of explanation involved is not the usual theory-laden notion that is subject to pragmatic, contextual parameters but is supposed to be some other metaphysical notion. For the sake of space, it will be best to avoid this issue and the notions of grounding that go with it. I will, however, assume that there is some useful notion of ontological dependence available to us.

prior value of delicious, aromatic coffee, the domain-internal value of the machine is derivative.[53]

Goldman and Olsson describe themselves as "process reliabilists" about knowledge, summarizing their view in the following way: "[S]ubject S knows that p if and only if (1) p is true, (2) S believes p to be true, (3) S's belief that p was produced through a reliable process, and (4) a suitable anti-Gettier clause is satisfied." (2009, pg. 22) Process reliabilism about knowledge, of course, has an illustrious history and much ink has been spilled (by Goldman in particular) in precisifying and defending the view. Here, Goldman and Olsson are content to assume that readers are at least somewhat familiar with the details of the view as they attempt to provide a solution to the value problem.[54] For them, knowledge is necessarily produced by a process that reliably yields true beliefs while mere true belief is not. That, in some sense, is supposed to make knowledge epistemically better.

Reliable mechanisms that produce valuable things are valuable things themselves. They are *instrumentally* valuable entities.[55] Within a particular critical domain, reliable mechanisms are *derivatively* valuable entities as the value of the mechanism *depends* on the prior value of the domain-internally good entities that it reliably produces. Given that Goldman and Olsson are veritists, one might imagine the following sort of "bottom-up"

---

[53] If one prefers a modal-existential reading of ontological dependence claims, what is asserted here is that, necessarily, a coffee machine that produces delicious, aromatic coffee has value within the "world of coffee" only if delicious, aromatic coffee has value within that "world." Alternatively, it is necessary that if delicious, aromatic coffee does not have coffee-related value, then neither does a machine that produces such coffee.

[54] For a more details on the process reliablist view, see Goldman (1979; 1986). Goldman thinks of reliability as constituting a justification condition on knowledge.

[55] Generally speaking, in this chapter, I will use the notion of instrumental value relatively liberally to include not only certain causal and probabilistic relations but also certain constitutive relations. For example, pleasure would still be instrumentally valuable relative to happiness even if it does not cause happiness or make happiness probable, but is rather a constitutive means to happiness in some other sense.

process reliabilist approach to the value problem: True beliefs are epistemically good as fundamental epistemic values. Knowledge, though, is better. Since knowledge is necessarily the product of a reliable mechanism, it always has an added derivative value.

Goldman and Olsson, however, are aware that this way of answering EVP cannot be satisfactory. As Linda Zagzebski (among others) has shown us, being the product of an instrumentally valuable mechanism does not add value to the product itself:

> The good of the product makes the reliability of the source that produced it good, but the reliability of the source does not then give the product an additional boost of value… If the espresso tastes good, it makes no difference if it comes from an unreliable machine… [By analogy,] if the belief is true, it makes no difference if it comes from an unreliable belief-producing source. (2003, pg. 13)[56]

The central idea here is that, even though *being a reliable producer of good things* is a valuable property for a mechanism to have (i.e., a derivatively valuable one), *being the product of such a mechanism* does not add any relevant value to that product. Being the product of a reliable espresso machine does not make a delicious and aromatic cup of espresso any better. Similarly, if we take true belief to be the relevant fundamental epistemic value, being the product of a reliable, truth-conducive process does not make some particular true belief any better. Even more, being the product of some reliable mechanism would not make a *bad* cup of espresso or a *false* belief any better either. Products simply do not "borrow lustre" from their sources, at least when those sources are only instrumentally valuable relative to the goodness of the products themselves.

---

[56] See also Zagzebski (1996; 2000). Jones (1997), Swinburne (1999), Riggs (2002), Kvanvig (2003a), and Sosa (2003) have made similar points about reliable sources and their products. Some, including Zagzebski, raise the issue in order to carry out an explicit attack on the process reliabilist account of knowledge, arguing that process reliabilists cannot thereby explain the greater value of knowledge over mere true belief. Since I have made the distinction between modal and essential formulations of the value problem, I do not take a particular analysis's failure to explain the added value of knowledge over mere true belief as a theoretical failure *if* some other non-essential, modal response to EVP can be given. Here, I will focus more directly on Zagzebski's point as it contrasts with Goldman and Olsson's actual purported solution to the value problem.

In the literature, this Zagzebskian point has been framed in a number of ways, sometimes in terms of the "swamping" of instrumental value.[57] Goldman and Olsson are sensitive to the "swamping" point and they attempt a different kind of bottom-up, process reliabilist solution to the value problem that will avoid such objections. Consider their "conditional probability solution:"

> [I]f a true belief is produced by a reliable process, the composite state of affairs has a certain property that would be missing if the same true belief weren't so produced. Moreover, this property is a valuable one to have – indeed, an epistemically valuable one. Therefore, *ceteris paribus*, knowing that *p* is more valuable than truly believing that *p*. What is this extra valuable property that distinguishes knowledge from true belief? It is the property of *making it likely* that one's future beliefs of a similar kind will also be true. More precisely, under reliabilism, the probability of having more true beliefs (of a similar kind) in the future is greater conditional on *S*'s *knowing* that *p* than conditional on *S*'s *merely truly believing* that *p*. (2009, pgs. 27-8)

Granting Goldman and Olsson the truth of process reliabilism for the sake of argument (which, of course, could certainly be challenged), the solution here is supposed to be able to avoid Zagzebskian worries about the failure of instrumentally valuable mechanisms in adding value to their products.[58] How? Goldman and Olsson move to talking about the "composite state of affairs" of a relevant agent's cognitive state and the way in which it was produced. For them, the "composite state of affairs" involved in gaining knowledge is epistemically better than what it would be otherwise since it includes both a good cognitive state (i.e., a true belief) *and* a mechanism that is epistemically good (i.e., one that is a reliable producer of true beliefs). I will come back to discuss this move to the "composite state of affairs" momentarily.

Goldman and Olsson illustrate their solution by means of an espresso example:

---

[57] See Kvanvig (2003a) for the "swamping" terminology.
[58] I will have more to say about the relationships between analyses and the value problem when I discuss Greco's virtue reliabilist solution later in this section.

> If a good cup of espresso is produced by a reliable espresso machine, and this
> machine remains at one's disposal, then the probability that one's next cup of
> espresso will be good is greater than the probability that the next cup of espresso
> will be good given that the first good cup was just luckily produced by an
> unreliable machine… [T]he reliable production of a good cup of espresso…
> raise[s] or enhance[s] the probability of a subsequent good cup of espresso. This
> probability enhancement is a valuable property to have. (pg. 28)

Their point is that what is involved in reliably producing a particular good cup of

espresso is the use of a reliable mechanism. Insofar as that mechanism is still available

for future use, the probability of having more good cups of espresso is greater than it

otherwise would be. The "composite state of affairs" of coming to reliably produce a

good cup of espresso also includes the possession of a good mechanism and that

mechanism is good because, by possessing it, you more likely to produce more good cups

of espresso across time.

Similarly, for Goldman and Olsson, what makes the composite state of affairs of

knowing that p better than the state of affairs of merely truly believing that p is that

coming to know includes a mechanism that makes it probable that we will have more true

beliefs in the future. The purported solution to the value problem here is a unique one but

it should remind us of the other cross-temporal attempts encountered in the first section

of this chapter. Recall that there are usually exceptions to such cross-temporal claims.

For instance, Williamson's claim about how well knowledge-constituting true beliefs are

retained across time is subject to the condition that one's cognitive faculties are "in good

order," a condition that might easily fail to obtain. Similarly, Goldman and Olsson are

quick to point out that their cross-temporal claims depend on a number of "empirical

regularities." According to them, in order for the purported probability enhancement to

accrue, the reliable mechanism at work must be a sufficiently general one, the agent must

have cross-temporal access to that very same mechanism, and she must be likely to use it again in similar future circumstances. (pgs. 29-30) By their own admission, these conditions will not always hold.

This should cause us to worry about Goldman and Olsson's purported solution as it falls short of the modal intuition assumed in EVP. What needs an explanation is not why knowledge, in cases where certain empirical conditions are met, is epistemically better than mere true belief. We need to explain why knowledge is *necessarily* epistemically better than mere true belief. Even in cases where Goldman and Olsson's empirical regularity conditions fail to obtain, knowledge is still better than mere true belief from the epistemic perspective. Take an example: Suppose that *S* comes to know that *p* through some reliable perceptual mechanism but also knows that the perceptual mechanism used will not be available to her for future use (perhaps *S* perceptually *sees* that *p* but also knows she is about to go blind). *S*'s epistemic situation vis-à-vis *p* is *still* better than what it otherwise would be if *S* had only luckily come to a true belief. Even when there is no cross-temporal access to the relevant knowledge-generating reliable mechanism, the token case of knowledge is still epistemically better than a corresponding token of mere true belief. Goldman and Olsson's "conditional probability solution" cannot account for this.

Goldman and Olsson are mostly unconcerned with accounting for a *modal* version of the value problem of knowledge: "[I]t is far from clear that our pre-systematic thinking on the matter demands that knowledge always be more valuable than mere true belief." (pg. 31) However, they simply do not make the distinctions required in order to clearly see which modal claims are suspect and which others are more genuine. I agree with

Goldman and Olsson that the modal assumptions involved in VP1, VP2, and predicative readings of EVP might easily be undermined but, given what has been claimed in the previous section, that does not make an attributive reading of EVP any less pressing. We should expect a real solution to an attributive reading of EVP and, hence, Goldman and Olsson's reliabilist solution does not fully capture the explanatory desiderata.

In order to avoid Zagzebskian worries about instrumental value, Goldman and Olsson move to a comparison of the overall "composite state of affairs" involved in gaining knowledge as opposed to coming to a mere true belief. This move is also suspect. When we ask about the epistemic superiority of knowledge over mere true belief, it is plausible to think that we mean to compare the *cognitive states themselves* and not merely some cognitive state/causal-production hybrid entities (i.e., the "overall" state). We think that knowledge, which is a kind of cognitive state, is (*ceteris paribus*) epistemically better than mere true belief, another kind of cognitive state. Perhaps the cognitive state of knowledge is better than the cognitive state of mere true belief *because* it has an essential kind of causal production but, plausibly, moving to compare some "composite states of affairs" is a red herring.

Perhaps, though, there is more to Goldman and Olsson's move. As Berit Brogaard (2006) demonstrates, a theorist about the value of knowledge adhering to the so-called "Moorean conception of value" might feel the need to shift from a comparison of the relevant cognitive states (i.e., knowledge and mere true belief) to a comparison of the overall situations involved in knowing and mere true believing. The Moorean conception of value is the following: If two entities x and y have the same intrinsic properties, then x and y are equally valuable. If one takes the Moorean conception to be

true and also thinks that knowledge is a special kind of true belief, then the only way to explain the superiority of knowing over mere true believing would be to move to a comparison of the broader situations. Since a given token of knowledge can only be individuated from a corresponding token of mere true belief via its extrinsic properties (e.g., its causal history, its relations to the agent's other cognitive states, the environment in which it is produced or maintained, etc.), a proponent of the Moorean conception must think that, *qua* cognitive state, knowledge and mere true belief are equally valuable. Hence, the only way to save the superiority of knowing would be to move to a comparison of the relevant overall "composite states of affairs."

Even though a value theorist might be so motivated, Brogaard draws on the work of Wlodek Rabinowicz and Toni Roennow-Rasmussen (1999; 2003) to show us that the Moorean conception of value is flawed. Intrinsic duplicates do occasionally generate differences in value. For example, the original Les Demoiselles d'Avignon is more valuable than an intrinsic duplicate simply because the former but not the latter was painted by Picasso. The painting's extra value depends on its extrinsic relation to Picasso.[59] Similarly, it seems possible that some bit of knowledge is more valuable than an intrinsically identical true belief in virtue of its extrinsic properties. This possibility of an added extrinsic value seems particularly relevant when we are speaking about domain-internal value (like we are in an attributive reading of EVP). For example, the original

---

[59] Notice that this does not run afoul of the Zagzebskian claim that no value can be added to a product merely by being produced by an instrumentally valuable source (relative to that product). It is not as if the real Les Demoiselles d'Avignon is more valuable than an intrinsic duplicate because being produced by Picasso is merely instrumentally valuable relative to paintings that look a particular way. After all, a good Picasso imitator would be just as instrumentally valuable with respect to *that*. The relevant value-theoretic relations are different than those accruing between, say, good cups of coffee and being produced by a reliable machine. The extrinsic relations accruing between the real Les Demoiselles d'Avignon and Picasso add a certain *fundamental* artistic value to the painting, not a mere derivative, instrumental value.

Les Demoiselles d'Avignon is more domain-internally valuable than some intrinsic duplicate of the painting from the perspective of an art collector. Within certain critical domains like art collecting, the extrinsic properties of things often make a world of difference. Why would it not be the case that something similar holds when we are discussing knowledge's epistemic value?

Consequently, Goldman and Olsson's move to discussing the differences in value between the relevant "composite states of affairs" does not seem warranted (at least not for reasons related to the Moorean conception). In EVP, we want to know what explains the difference in epistemic value between the *cognitive states* of knowledge and mere true belief. This is an important point because it limits what can plausibly count as a good bottom-up solution to the value problem. If true belief has fundamental epistemic value and we are directly comparing the cognitive states of knowledge and mere true belief, then the fact that knowledge is produced in a way that is *instrumentally* valuable relative to true belief cannot possibly account for knowledge's epistemic superiority. The value of the product cognitive state will always be "swamped" by whatever value (or disvalue) it has otherwise as constituting a true (or false) belief.

This is not merely a problem for purported reliabilist solutions like Goldman and Olsson's. It is a problem for any bottom-up theory about the value problem that (1) purports to explain the added epistemic value of knowledge in terms of the value of some justificatory status and (2) takes that justificatory status to be merely *instrumentally* valuable relative to the epistemic value of true belief. For example, even an access internalist bottom-upper will not have avoided Zagzebski's worries if she takes her internalist justification to be merely instrumentally valuable relative to the value of true

belief. Suppose that she takes being access internalist justified to be epistemically valuable because it makes a given belief objectively more likely to be true. Even if it is the case that believing justifiably does enhance the probability of a belief's being true, such a theorist would run headlong back into worries about the swamping of the value of the resultant cognitive states. Even if S comes to a true belief that p via a procedure that meets the strictures of access internalist justification, given that justification is merely instrumental to obtaining true beliefs, that fact cannot possibly *add* epistemic value to S's belief that p over and above the value it already has in virtue of being true.[60]

Because we are comparing the epistemic value of the cognitive states of knowledge and mere true belief and not some broader hybrid entities, *instrumentalist* bottom-up solutions seem bound for failure.[61] Kurt Sylvan (forthcoming) has recently become aware of this, attempting to develop a non-instrumentalist solution to the value problem. Like Goldman and Olsson, Sylvan is a veritist and he hopes to explain knowledge's epistemic superiority to mere true belief in terms of the *derivative* value of other essential properties of knowledge. What makes Sylvan's attempted solution unique is that he refrains from using the purported *instrumental* value of properties like epistemic justification to explain knowledge's superior status.

---

[60] See Pritchard (2010, pgs. 9-11).

[61] Goldman and Olsson do put forth another distinct solution to the value problem that is not an instrumentalist one. They call it their "value autonomization" solution (see pgs. 31-5). I avoid talking about this solution in detail for a number of reasons. First, the value autonomization solution, like the conditional probability solution, seems to rely on Goldman and Olsson's problematic move to comparing the "composite states of affairs" involved in knowing and mere true believing and it is subject to the same worries that I have brought out above. Second, given the nature of the solution and its attribution of "autonomous" epistemic value to items other than mere true belief, it is unclear as to whether it is a proper bottom-up solution in the way that I have described and it is highly questionable as to whether it is even a genuine veritist solution, which is surprising given Goldman's other commitments. Third, my most general worry about bottom-up solutions will apply equally to their view on value autonomization. I discuss this worry at the end of this section.

Sylvan rejects what he calls "instrumentalism" about derivative value relations within critical domains. Instrumentalism is the following view: If x has derivative value within domain D, then x either produces entities that are fundamentally valuable within D or x is the product of something that produces such entities. For an instrumentalist, derivative value is always grounded in the obtaining of certain kinds of instrumental relations. Again, instrumental value is *derivative* because it asymmetrically depends on the value of the more fundamental entities. By denying instrumentalism, Sylvan leaves open the possibility that derivative value might accrue to entities in virtue of their entering into other *non-instrumental relations* with more fundamentally valuable entities.

He uses a case of moral evaluation in order to illustrate his non-instrumentalist stance. According to Sylvan, beneficent actions are morally valuable and, yet, so are actions that "manifest a proper valuing of beneficence" even if such actions do not themselves end up being beneficent.[62] Suppose that Alice intends to help a person in need by performing some action but, unluckily, her action causes more harm than good. Alice's action, even though it did not turn out to benefit anyone and was not technically a beneficent one, still has moral value. Even though the action does not manifest beneficence, it nonetheless manifests a proper way of valuing beneficence and that is enough to make it worthy of some moral goodness.[63] According to Sylvan, the moral value accruing to Alice's action is derivative but non-instrumental. He is a proponent of a principle that he attributes to Thomas Hurka: "When V is a non-instrumental value from

---

[62] Sylvan, here, seems to interpret beneficent action as *successful* beneficence, actually bringing about some good for an intended beneficiary.

[63] An ethicist might easily interject here and claim that Alice's action itself is not the proper item of moral worth. It is Alice's *intention* which has the relevant moral value. Though this would be an entirely reasonable way of pressing Sylvan's analogy (and, more generally, Hurka's principle described below), I will not focus on it here and will, for the sake of argument, assume that actions can take on moral value *in virtue* of their associated intentions.

the point of view from domain *D*, proper ways of valuing *V* in *D* and their manifestations have some derivative non-instrumental value in *D*." (pg. 5) For the case at hand, suppose that beneficent action has non-instrumental moral value. By Hurka's principle, proper ways of valuing beneficence and their manifestations (including Alice's action) would have derivative but non-instrumental moral value relative to the value of beneficence.

The moral value accruing to Alice's action is *derivative* relative to the moral value of beneficence because the value of Alice's action asymmetrically depends on the moral value of beneficence. If beneficence were not morally valuable, then her action would not be either. The moral value accruing to Alice's action is also *non-instrumental*. The action's derivative value is not grounded in its instrumental relations to beneficence. Even if Alice inhabited a world where some demon secretly but consistently thwarted any attempt at beneficent action, Alice's action would seemingly still have moral value in virtue of its manifesting a proper way of valuing beneficence. Acting in the way that Alice did could not be instrumentally valuable relative to the value of charity in that world because it would never actually bring about good results. Nonetheless, such actions could still have moral worth.

Sylvan thinks that such non-instrumental derivative value relations are important because they might assist the veritist in avoiding worries about "swamping." According to Sylvan, the accruing of non-instrumental derivative value relations, unlike instrumental relations, *can* add value to a product that has a more fundamental domain-internal value:

> [S]uppose 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 [moral] 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. (pg. 5)

In the example, even though Alice's action is beneficent and has moral value in virtue of that fact, there is an *extra* value that accrues to her action (and not Beatrice's) because it also manifests a proper way of valuing beneficence. This is true even though the extra value accruing to Alice's action is derivative relative to the moral value of beneficence. Hence, it seems to Sylvan that derivative value relations occasionally can add value to certain fundamentally valuable entities so long as those relations have the right non-instrumental form.

If Sylvan could demonstrate that the comparison between knowledge and mere true belief has an analogous structure to that of Alice and Beatrice's actions (i.e., in the sense that the former manifests a proper way of valuing some more fundamental domain-internal good), then perhaps he and other veritists would have a good answer to EVP. Purportedly, if knowledge manifests a proper way of valuing true belief, then it could have an added derivative value over and above mere true belief, even if knowledge just is TB+. A theorist about epistemic value could then hold on to veritism while also providing a satisfactory explanation of knowledge's superior epistemic status.

This is Sylvan's strategy. Again, the particular derivative value relation that he focuses on is that of *properly valuing* something good within a domain. There are many ways for a subject S to properly value a class of good things G: S might be dedicated to G, protect G, take delight in G, respect G, comply with G as an ideal, etc. Not all these ways of valuing will necessarily instrumentally promote G. Yet, insofar as they are proper ways of valuing G, then the obtaining of such relations will generate derivative value that can add value to a product. For example, Alice's action is beneficent and is

thus morally good. But, her action also manifests a *respect* for beneficence and, hence, it has an added derivative value. Sylvan's idea is that perhaps other intentional states like belief can be thought of in similar ways.

Suppose that Alice comes to a true belief that *p*. For the veritist, true beliefs have fundamental epistemic value. Now also suppose also that Alice's true belief that *p* constitutes knowledge. Why would that be any better than Alice *merely* coming to a true belief that *p*? Sylvan's answer: It would be better because Alice's belief would also manifest a "compliance" with the norm of true belief (i.e., being correct to believe that *p* iff *p* is true). For Sylvan, compliance with a relevant norm is a proper way of valuing some good thing because it occurs when one conforms to the norm by "strongly respecting" it (in the sense that one's conformity is guided by objectively good reasons). Hence, since Alice's belief is a proper way of valuing truth/accuracy in this way, it would have an added derivative value. Purportedly knowledge essentially manifests such a feature; by its very nature, it is a state that manifests a conformity with the norm of truth (or accuracy) *via* strong respect.

Consider a crude example of the view: Suppose that Alice heeds the proper evidence for and against *p* and comes to believe that *p* on the basis of it. Insofar as Alice then *knows* that *p,* it is plausible to think that her belief conforms to the norm of truth by manifesting a strong respect for it. After all, the belief will be justified and that justification will in some sense be responsible for the belief's accuracy and its conformity to the norm. *Strongly respecting* the truth norm in this way at least initially seems to be a proper way of valuing the truth and it can then purportedly account for knowledge's added value.

For the sake of argument, let us grant to Sylvan that properly valuing a good thing is a legitimate, non-instrumental way for derivative value to be added to a product. Nevertheless, in order for such an approach to meet the requirements of EVP, it must satisfy the modal intuition about the value of knowledge. Knowledge is *necessarily* epistemically better than mere true belief. Is it really the case that knowledge everywhere and always manifests a proper valuing of accuracy as Sylvan imagines? Sylvan does consider the worry that his account is too demanding: "[M]ust we place value on accuracy in thought to believe… justifiedly or knowingly?" (pg. 32) Consider an example: Adam randomly and quite accidentally glances to the left side of his desk and he sees that there are six pencils laying there. Adam automatically forms the belief that there are six pencils there and, hence, comes to know that there are six pencils there. Does Adam's knowledge necessarily manifest a way of *valuing* truth or accuracy with respect to whether there are six pencils there? Does the fact that his belief is justified necessarily manifest a *respect* for the truth about that issue? It seems not. It seems perfectly possible that Adam does not care one bit about the accuracy his belief and, yet, still believes justifiably and knowingly.[64]

Sylvan responds to such worries about demandingness:

> This objection itself is nourished by overly demanding assumptions. In particular, it assumes that placing value on accuracy in thought is more demanding than it is… When we reason carefully and heed the apparent evidence bearing on whether we would conform to a norm N by ø-ing, it is plausible that our ø-ing exhibits respect for N… We need not view respect [and certain other ways of valuing]… as constituted by emotions or states with any fancy qualitative feel.

---

[64] I take it that the following considerations also raise similar worries for Zagzebski's (2003) solution to the value problem. For her, since mere reliability and instrumentality cannot solve the value problem, the extra value of knowledge over mere true belief must be accounted for in a way that is analogous to how intentions add value to actions. In fact, by her account, knowledge *just is* a kind of action. The problems I raise for Sylvan's view should apply *mutatis mutandis* to her view as well.

They are just ways to be attuned to different factors that bear on the accuracy of particular beliefs. (pg. 32)

We are invited to think that it does not take much to place value on accuracy or truth in believing. Indeed, Sylvan construes what it takes to value something very broadly, including "any way of being positively oriented toward something in attitude, act, or disposition." (pg. 4) In coming to believe that *p*, we need only to have unconsciously and automatically activated the appropriate *dispositional capacities* in order to count as having manifested the right kind of respect for truth. Purportedly, Adam manifests a strong respect for accuracy in his believing solely in virtue of his automatic activation of those truth-tracking dispositions, regardless of whether or not he has any positive pro-attitudes or desires about the accuracy of his belief.

It seems fair to question whether we are working with the right notions of valuing and respect here. We can further specify the example that we are considering to make the relevant issue clearer. Suppose that Adam does not merely lack a desire about accuracy vis-à-vis his beliefs about the items that are on his desk but, instead, he actively despises following an accuracy related norm with respect to his beliefs about such things. Perhaps he feels bad about how often his desk is cluttered and, instead of simply cleaning his desk, he actively works toward believing that it is clean and empty, regardless of whether that is true. Yet, sometimes, his automatic processing and dispositional capacities get the best of him and he forms the relevant accurate beliefs about the number of pencils he has laying around his desk. When he forms such beliefs on the basis of his perceptual experience, he is not merely accurate, he also forms his beliefs justifiably and in a way that they constitute knowledge, even though he forms such beliefs against his own reflective wishes. When he forms such beliefs, is Adam properly *valuing* accuracy?

Let us suppose for the sake of argument that there is a legitimate sense in which Adam's belief does manifest a respect for accuracy (i.e., Sylvan's "strong respect"). A question can now be raised about whether manifesting that kind of respect is really something that can count as a proper valuing of truth in a way that will *add* value to the resultant cognitive state. Take another ethics analogy: Suppose that opening doors for the elderly is a morally good thing and that Brian occasionally does so. Now, suppose that we also learn the following about Brian: He has been raised as a so-called "southern gentleman," has acquired an unreflective and automatic disposition to open doors for the elderly, but he also despises his upbringing and hates the disposition that he has. However, Brian's door-opening dispositions occasionally activate against his more reflective wishes and he ends up opening doors for the reasons he was taught as a child. Given Sylvan's requirements, it would seemingly be appropriate to say that Brian's actions occasionally manifest a relevant kind of respect for a moral norm about door-opening solely in virtue of the automatic activation of his door-opening dispositions.

Nevertheless, I doubt that we would be tempted to say that a token door opening of Brian's is *morally better* than it otherwise would be by such activation. This is because Brian's situation is relevantly different than the more uncontroversial cases of derivative value addition considered earlier. As Sylvan states, we think that Alice's action is morally better than Beatrice's because it manifests a genuine respect for beneficent action. What really makes her action better, however, is that she has a conscious *intention* to do good for goodness' sake, unlike Beatrice. She has a conscious *purpose* to benefit someone else. That is what makes Alice's respect of beneficence the kind of valuing that can *add* moral value to her action. Brian's case is relevantly

different. Even if there is some sense in which Brian manifests a respect for moral goodness in the mere activation of his dispositions, it is a less robust sense of respect that should make us wonder whether any value has been added.[65] Similarly, we should wonder whether the activation of Adam's truth-tracking dispositional capacities really manifests a respect for truth in a way that can add epistemic value to his resultant belief.

Ultimately, this is to question whether Sylvan has really given us a genuine answer to EVP. Even though non-instrumental derivative value relations within critical domains are important to consider, we should be wary that the particular relations that Sylvan points out can do the trick needed for a satisfactory veritist, bottom-up solution to the value problem. This is because there seem to be many cases of knowledge in which the knowing subjects have no positive pro-attitudes or desires about truth and accuracy. This is not to claim that positive instances of valuing require a "fancy feel," but it is to say that the right kind of respect requires *some* positive qualitative feelings (or at least a disposition to have such positive qualitative feelings) toward some object. Knowers simply need not manifest this kind of respect for truth, at least when the sense of respect in question more uncontroversially adds value to a cognitive state like it does to an action vis-à-vis moral evaluation. Yet, even in cases where such respect is absent, knowledge is still epistemically better than mere true belief.

If the epistemically valuable properties of knowledge, those derivative to the value mere true belief, are going to be able to provide us with the necessary resources for

---

[65] To be clear, the point here is not that Brian's *disposition* to open doors for the elderly does not have any kind of moral value. The point is that, even if the disposition does have derivative value, it is not at all obvious that the value accruing to the disposition is then transferred to the action in the way that Sylvan imagines it does with Alice. In terms of value addition, the disposition in Brian's case seems more closely related to items that take on a derivative, instrumental value than those kinds of valuings that uncontroversially add value to a product.

a bottom-up solution to the value problem, they cannot be merely instrumental to true belief (as they are in Goldman and Olsson) and they cannot require too much of an agent with respect to her pro-attitudes. John Greco's (2010, Ch. 6) virtue reliabilist solution is worth discussing because it at least initially seems to avoid both of these obstacles. One issue in discussing Greco is that he, like many others, does not make the distinctions necessary to have a clear view of the landscape involved in the value problem (e.g., he does not make the distinction between genuine value and domain-internal value). Nonetheless, I still think that there is a way to situate his view within the framework I have laid out as giving a possible answer to EVP.

Greco argues that knowledge manifests an *achievement* of true belief in the sense that, in knowing, an agent *deserves credit* for her true belief: "(KSA) *S* knows that *p* if and only if *S* believes the truth (with respect to *p*) because *S*'s belief that *p* is produced by intellectual ability." (pg. 71) According to Greco, knowledge is an intellectually virtuous true belief where the exercise of cognitive ability and virtue figures importantly in a causal explanation of the belief's accuracy. Suppose, for example, that *S* truly believes that the square root of 2 is irrational. *S knows* that the square root of 2 is irrational only insofar as the accuracy of *S*'s belief is causally explained by the exercise of her relevant cognitive ability. *S* can know that the square root of 2 is irrational if she uses her inferential abilities, deriving it from the premises in Euclid's proof but not if she merely has a lucky hunch. Insofar as she knows, *S* deserves credit for her true belief because, in coming to believe, she exercises her cognitive abilities in a causally efficacious way. Greco argues that KSA offers us the correct verdicts in a wide range of other cases,

including those of perception, testimony, innate knowledge, and a variety of Gettier scenarios.[66]

He also believes that KSA can properly account for the value problem. Supposing that *x* is a valuable state of affairs, it is better to *achieve x* than to merely have *x* obtain luckily and by accident. An accurate belief is an epistemically good thing to have, but *achieving* accurate belief is better. Because knowledge, by its very nature, is supposed to be an agent's achievement of true belief, virtue reliabilists like Greco claim to have an answer to EVP. Consider, by analogy, the evaluation of an archery shot. A shot is better from the perspective of archery when it is accurate as a result of ability than when it is accurate as a result of luck.[67] The Olympian archer's accurate shot is better than the shot by Joe Schmo who luckily hits the bulls-eye despite his terrible form. In the case of the Olympian, the accuracy of the shot is, in some sense, attributable to the archer where it is not so attributable to Joe Schmo. Similarly, a belief is always epistemically better when it is accurate as a result of the agent's relevant cognitive ability (i.e., when it constitutes knowledge) than when it is accurate as a result of luck. The agent deserves credit because she has achieved accuracy.

This particular virtue reliabilist solution is another bottom-up approach, but perhaps less obviously so than those previously encountered. By Greco's account, certain beliefs have an added epistemic value because they constitute an achievement of true belief, being accurate because of the exercise of an ability. Constituting an

---

[66] See, in particular, Greco (2003; 2010, Ch. 5).

[67] The archer example is used, most famously, in Sosa's (2007) virtue reliabilist account of knowledge. There are some differences in detail between Sosa and Greco's accounts but the archer example works well for the purposes here. When there are differences between Sosa and Greco's accounts that make a difference with respect to the value problem, I will note them.

achievement of true belief, however, is *derivatively valuable* relative to the value of true belief. In the same way that the value of respecting a good thing depends on the value of that good thing (*a la* Sylvan), the value of achieving something also depends on the value of that thing. No one, for example, would be tempted to think that achieving world poverty is good. World poverty is not a valuable state of affairs and, hence, neither is the achievement of such. In fact, it would be tempting to say that some state of affairs cannot even count as a legitimate achievement unless the object of purported achievement itself has value. Either way, the value of achievements exhibits an asymmetric ontological dependence on the value of the state of affairs being achieved. Like other bottom-up solutions, then, knowledge purportedly has a derivatively valuable property relative to the value of true belief that makes it better than mere true belief.[68]

Yet, unlike other bottom-up solutions, Greco's approach at least initially seems to avoid the objections brought out above. Firstly, the derivatively valuable property under discussion (i.e., constituting an achievement of true belief) is not merely *instrumental* relative to true belief and thus will not be subject to worries about "swamping." In fact, Greco takes the Aristotelian line that goods achieved through the exercise of one's abilities or virtues are intrinsically valuable and are constitutive of human flourishing.[69] (2010, pg. 97) Secondly, achieving true belief through the exercise of one's cognitive

---

[68] While Greco's (2010) purported solution will count as a bottom-up solution, he does not commit himself to veritism like Goldman, Olsson, and Sylvan. In fact, there are certain passages suggesting that Greco is a pluralist about fundamental epistemic value (see, e.g., pgs. 7-8).

[69] It is not completely clear that being constitutive of human flourishing will necessarily contribute any extra *epistemic* value in the attributive sense to a cognitive state. If Greco's view is correct, constituting such an achievement would most obviously contribute to eudaimonia which is naturally associated with a kind of genuine value. However, to be most charitable, I will here assume that he is successful in showing us how a cognitive state's epistemic value can be enhanced by constituting an achievement. After all, it does seem to be the case that "success because of ability" contributes an extra archery-related value to an Olympian's shot as compared to Joe Schmo's independently of whether such shots contribute to eudaimonia.

abilities does not seem to require too much in the way of an agent's pro-attitudes. Even if

a knower has no reflective desires at all for accuracy in believing, that agent's cognitive

abilities may still be invoked in a relevant causal explanation of why her beliefs turn out

to be accurate. In that way, such cases of knowing will still count as achievements for

Greco, generating a greater epistemic worth.[70]

Like any account of knowing, however, it is not entirely clear that Greco does in

fact give us an adequate analysis of knowledge. Although Greco's virtue-theoretic

solution to EVP would avoid the specific pitfalls discussed in considering Goldman and

Olsson and Sylvan's solutions, I would like to show why this uncertainty in analysis

matters, generally speaking, in the way that we theorize about the value problem. In his

analysis, Greco writes off most Gettier scenarios as cases of deviant causal chains in

which the relevant cognitive abilities do not figure properly into a causal explanation of

accuracy, his approach to Goldman-style (1976) barn façade cases seems particularly

problematic. Consider Fake Barn Country:

> Henry is driving in the countryside and sees a barn ahead in clear view. On this
> basis he believes that the object he sees is a barn. Unknown to Henry, however,
> the area is dotted with barn façades that are indistinguishable from real barns from
> the road. However, Henry happens to be looking at the one real barn in the area.

Many epistemologists have had the intuition that, despite Henry's justified true belief that

there is a barn there, Henry does not *know* that there is a barn there because his accurate

belief is lucky in a relevant way. Given a commitment to KSA, however, Greco seems

---

[70] Duncan Pritchard (2010) has questioned whether being a success because of an ability is sufficient
enough to be counted as an achievement in a way that will add value to a result. He worries that some
successes because of ability do not merit such a status because they do not involve the overcoming of a
significant obstacle or (as I have pointed out here) need not exhibit any conscious motivational states. (pg.
29) Though I will discuss some of Pritchard's worries in detail below, I will not focus on this particular
point because it is not clear to me that his points apply to using Greco's approach in answering EVP.
Pritchard's point about "easy achievements" seems more aimed at the addition of *genuine* value, not
necessarily domain-internal value.

committed to the opposite verdict. After all, Henry has an accurate belief that there is a

barn there and a seemingly relevant cognitive ability (e.g., his perceptual belief-forming

ability) plays a significant role in a causal explanation of that accuracy.

Greco, however, wishes not to flout the traditional intuition about barn façade

cases. Consider his response:

> First, we note that abilities in general are always relative to environments. For
> example, Derek Jeter has the ability to hit fastballs relative to normal
> environments for playing baseball. He does not have that ability relative to an
> active war zone where he would be too distracted to focus on the ball. In general,
> when we attribute an ability we have in mind some relevant environment as well
> as relevant conditions, etc. Second, we may now claim that Henry does not have
> the ability to tell barns from non-barns relative to the environment he is in.
> Relative to normal environments, we may assume that Henry can perfectly well
> discriminate between barns and non-barns. Relative to Fake Barn Country,
> however, Henry does not have that ability. (2010, pg. 76)

The line here is that the possession of abilities, those that might enter into the relevant

causal explanations, is an environment-relative phenomenon. For Greco, in the same

way that Derek Jeter does not possess the ability to hit a baseball in an active war zone,

Henry simply does not possess the relevant knowledge-engendering ability in Fake Barn

Country (i.e., the ability to "discriminate between barns and non-barns"). Hence, since it

is not even possessed at the time, the relevant ability cannot be part of a causal

explanation of the accuracy of his belief. The intuition that Henry does not know is

allegedly preserved.

Duncan Pritchard (2010), however, pushes back against this way of thinking

about the barn façade case by introducing what he takes to be an analogous situation

involving an archer:

> [C]onsider a… case, in which Archie… selects a target at random, skillfully fires
> at this target, and successfully hits it because of his skill… Suppose, however,
> that unbeknownst to Archie there is a force-field around each of the other targets

such that, had he aimed at one of these, he would have missed. It is thus a matter of luck that he is successful, in the sense that he could very easily have not been successful. Notice, however, that luck of this sort does not seem to undermine the thesis that Archie's success is a genuine achievement… It is, after all, because of his skill that he is successful, even though he could very easily have not been successful in this case. That is, his success in this case is still primarily creditable to his archery abilities, even despite the luck involved in that success. (pg. 35)

According to Pritchard, Henry's case is similar in all relevant respects. Even though the environment that Henry is in makes it the case such that he might have easily been mistaken about whether there is a barn there, Henry still seems to possess and exercise the very same ability that he exercises in more normal environments and still seems to have properly *achieved* true belief in Greco's sense. For Pritchard, this undermines Greco's analysis of knowledge.

I take it that Pritchard's most pressing concern is the following: Even granting Greco's point about the environmental relativity of abilities, we need to be able to make a distinction between (1) an agent's lacking an ability because of the unfriendly environment she is in (e.g., Jeter's fastball hitting ability relative to an active war zone) and (2) an agent's being lucky that she is able to successfully exercise an ability in an otherwise unfriendly environment. Pritchard calls instances of (2) cases of "environmental luck" and Archie and Henry purportedly fit the mold. According to Pritchard, if we make the possession and exercise of abilities *so* context-sensitive that the presence of environmental luck automatically rules out the possession and exercise of an ability, then we seem to lose the distinction between (1) and (2) and thereby lose something important in explaining the nature of ability possession. In Pritchard's eyes,

this causes trouble for Greco's account of knowledge. Achievement seems compatible with the presence of a certain kind of environmental luck while knowledge is not.[71]

Now, even though I happen to side with Pritchard in this dispute, I am not raising the issue in order to take a hard line or say something final about it.[72] I bring it up to make the following point: Whether or not bottom-up solutions like Greco's are ultimately going to be successful will depend on the particular analysis of knowledge in question. With such bottom-up approaches, theorizing about epistemic value becomes handcuffed to metaphysical theorizing in a potentially problematic way. In Greco's case, his account of the epistemic value of knowledge stands or falls with the success of his achievement analysis of knowledge. Perhaps he is content with that fact but I would like to suggest that we should not be.

Suppose that Pritchard has successfully generated a class of genuine counterexamples to Greco's achievement model. The particular class that we were just considering would show that Greco's account is too weak, allowing for knowledge attribution when no such attribution is warranted (i.e., in certain cases involving environmental luck). If Pritchard's complaints are legitimate, then Greco's analysis does not pick out knowledge but only a kind of true belief that has a lesser epistemic standing

---

[71] One way around this particular difficulty would be to deny, as Sosa (2007) does, that environmental luck is incompatible with knowledge. This move, of course, raises different questions about the adequacy of the virtue reliabilist analysis. Whether or not it is a legitimate move, the points made below will still hold.

[72] To be fair, Greco is aware of Pritchard's Archie case and he does attempt (unsuccessfully, I think) to make a principled distinction between Archie and Henry (see Greco 2010, pgs. 88-9). Also see Kelp (forthcoming) for another line of criticism against virtue approaches like Greco's. An independent issue with Greco's line is that he consistently describes the relevant ability involved in Henry's case as that of "discriminating barns from non-barns" instead of the potentially more general way of describing the ability at work (e.g., that of believing on the basis of perception). Unsurprisingly, as Greco's approach is a version of reliabilism, there is a familiar difficulty here in identifying the correct level of description for the relevant abilities involved. This has been called the "generality problem" for reliabilism (see, e.g., Conee and Feldman 1998).

than knowledge does. It picks out a cognitive state that sometimes "falls short" of

knowledge. Hence, any appeal to achievement could only show us why *that* state is

epistemically superior to mere true belief, not knowledge. Even if achievement were,

contra Greco, merely a *necessary* condition on knowledge, that still would not distinguish

knowledge's epistemic value from the value of certain kinds of states that fall short of it

(although it would distinguish knowledge's value from the epistemic value of mere true

belief).[73] This, however, is part of what we want when we theorize about the value

problem. We want to know why knowledge has a *unique* standing as the ideal epistemic

relation obtaining between an agent and a true proposition.[74]

Again, we are theorizing about epistemic value in EVP. This handcuffing of

domain-internal value to a particular metaphysical view seems problematic for the

following reason: It seems that, *independently* of whether any particular analysis of

knowledge is correct, *independently* of whether philosophers like Greco have successful

responses to purported counterexamples, it will still be the case that knowledge has a

unique sort of epistemic standing. Ideally, we want *know* whether *p* and not have a mere

true belief about whether *p*, at least insofar as we have epistemic interests. Knowledge

has the unique epistemic standing that it does no matter what the correct analysis of it

turns out to be and, importantly, even if there is no reductive analysis of it to be given.

The bottom-upper, however, thinks of things differently. According to the bottom-upper,

what we are *really* after, most fundamentally, is true belief. Hence, in order to see why

---

[73] Pritchard (2010, pgs. 40-3) also argues against the necessity of the achievement condition by appealing to certain cases of testimonial knowledge. In any case, even if he does not successfully undermine Greco's achievement as a necessary condition on knowledge, he still takes himself to have shown that Greco's purported solution flouts "the secondary value problem" (even if it accounts for a difference in value between knowledge and mere true belief). The point I am making here is modelled after Pritchard's in that respect.

[74] See footnote 26 in Chapter 1.

knowledge is always epistemically better, we need to look at other aspects of its nature and search for helpful derivatively valuable properties. As shown above, this can be a heavy theoretical burden. It is not just a burden for Greco; it is a burden for anyone taking a bottom-up approach. In the final section, I will outline an alternative way of approaching the value of knowledge that can better accommodate these insights.

The final worry that I will bring out here is another general one, applying to all bottom-up solutions. Each of the bottom-uppers that we have encountered take it as a given that true belief has fundamental value within the epistemological critical domain.[75] In fact, veritists like Goldman, Olsson, and Sylvan take true belief to be the *only* kind of item that has such fundamental value. Many philosophers take true belief to lay at the bedrock of theorizing about epistemic value. Hilary Kornblith claims that "any account of epistemic evaluation that does not give truth a central role to play is inaccurate." (1985, pg. 158) Speaking about truth from the epistemic point of view, William Alston states: "[H]ow could any property of a belief be better…?" (2005, pg. 40) Even thoroughgoing pluralists like Michael DePaul think that the fundamental epistemic value of true belief is "uncontroversial." (2001, pg. 177) Allan Hazlett claims that "it is not coherent to ask whether true belief has epistemic value." (2013, pg. 262)

At the very least, I do think that the *fundamental* epistemic value of true belief is coherently questionable and, even more, it is questionable in a way that should make us doubt whether it ought to think of it as having such a central standing in epistemology. Compare the following two cases: (1) Jimbo occasionally forms his beliefs on the basis of

---

[75] In Greco's case, we cannot be sure that he thinks that true belief has *fundamental* epistemic value as he simply does not use the fundamental/derivative value distinction in drawing out his solution. Given the distinction that I have drawn here, however, we can be sure that the value of true belief is at least *more* fundamental than the other relevant epistemic desiderata in play for Greco in answering EVP.

tarot card readings. On the basis of a particular reading, Jimbo forms the belief that his

car is being towed. His car is not being towed. (2) Jimbo occasionally forms his beliefs

on the basis of tarot card readings. On the basis of a particular reading, Jimbo forms the

belief that his car is being towed. As things turn out, his car actually is being towed.

Jimbo's belief is false in case (1) and true in case (2). A relevant question is the

following: Restricting ourselves to epistemic desiderata, is Jimbo's situation any better in

case (2) than it is in case (1)? My own inclination is to say "no." Jimbo's beliefs seem

equally bad from the epistemic point of view. The fact that Jimbo's belief corresponds to

reality in (2) has no bearing on its epistemic worth.[76] This, I think, is something that the

bottom-up theorist will have a tough time explaining. If true belief really has

fundamental value within the epistemological critical domain, then why do we have such

judgments about the comparative epistemic value obtaining in (1) and (2)? Seemingly, if

the value of true belief is fundamental in a way that the other relevant epistemic values

depend on it, those other derivative values may be important, but true belief should still

receive at least a significant spotlight in epistemic value considerations. This seems to

indicate a sort of cognitive dissonance between our ordinary judgments about epistemic

worth (e.g., Jimbo's case) and thinking about true belief as having the kind of

fundamental epistemic value that the bottom-upper attributes to it.

To be clear, I am not claiming that Jimbo's belief in (2) is no better, in any sense,

than his belief in (1). His true belief in (2) could very easily be more *genuinely* valuable

than his belief in (1). It certainly seems *practically* better. After all, the second belief is

---

[76] This is to agree with Feldman (2002): "Imagine a person who makes an unreasonable and unreliable
inference that happens to lead to a true belief on a particular occasion. It might be fortunate that he's got
this true belief, but I see nothing epistemologically meritorious about it." (pg. 379)

accurate while the first one is not. What I am suggesting, however, is that such considerations do not necessarily bear on the epistemic point of view. This, of course, is not a bottom-up way of thinking about Jimbo's case. Consider what Sylvan might say: "Jimbo's belief in (2) has some epistemic worth, while it does not in (1). In fact, his belief in (2) manifests the only fundamental epistemic good, that of true belief. But, of course, it could be epistemically *better* since it does not manifest a *strong respect* for the truth." Or Greco: "Yes, Jimbo's belief in (2) has some epistemic worth, while it does not in (1). But, of course, his belief in (2) could be epistemically *better* since it does not manifest an *achievement* of true belief." My own inclination is to say that, from the epistemic perspective, we just do not really value certain kinds of true beliefs. The fact that a belief is accurate, by itself, does not seem to generate epistemic value, much less epistemic value of a fundamental kind.[77]

For all of these reasons, I think that it is questionable as to whether there is a satisfactory bottom-up solution to the value problem. Since we have only surveyed a number of considerations with respect to a few recent bottom-up approaches, this is not to be thought of as the last word on the matter. However, the points made should cause us to take a skeptical eye toward such answers to EVP and, more generally, to views of

---

[77] It is not completely clear as to whether this criticism can be applied to Sosa's (2007) virtue reliabilist view. While Greco's account of knowledge as an *achievement* of true belief assumes the antecedent value of true belief, Sosa's account of knowledge as *apt* belief might not make such an assumption. Apt belief is belief that is accurate *because* of a certain kind of exhibited adroitness (see Lecture 2). There are places in which Sosa seems to indicate that it is *not* mere true or accurate belief that we are most fundamentally after in our epistemic evaluations, but instead aptness. He likens the epistemic assessment of an apt belief to the assessment of a ballerina: "We had paid to see a performance, the product of artistic excellence and control… [C]ausation plays a crucial role. It's the movements under the dancer's *control* that we admire. But it is not clear that causation is there adding value to something that might still have existed but would not have had that value without its causal liaison." (pg. 76) Because of this, there may be ways of reading Sosa that would be in line with the view suggested here. Then again, there are other places in which he seems to suggest that mere true belief does have a fundamental value (see, e.g., pgs. 87-8). It is, of course, another question altogether as to whether the aptness analysis correctly captures knowledge.

epistemic evaluation that begin with the fundamental value of true belief. In our

epistemic evaluative practice, mere true belief just does not seem to play that kind of

fundamental role. I end this section with a quote from Alan Millar (2010) who expresses

this sentiment forcefully, but nicely:

> Since it is internal to the interest climbers have in grasping a rope that the grasp
> should be secure, it would clearly be absurd to try to account for the value of their
> having a secure grasp in terms of the value of a grasp on the rope that is barely a
> grasp. Knowledge is a mode of grasping the truth, and it would be no less absurd
> to try to explain the value of the grasp of the truth in which knowledge consists in
> terms of the grasp (if that is the right word) in which a merely true belief consists.
> (pg. 167)

2.3: The Goals of Inquiry and the Value Problem

If bottom-up approaches to the value problem do not satisfactorily meet our

explanatory needs, then what are some viable alternatives? If there are not any better

options on the table, then we may as well continue down the same path. After all, I think

we have seen at least some improvement on the bottom-up front and maybe that is some

reason to be optimistic that acceptable answers will eventually be available. However,

given the considerations above, if there are plausible alternative routes to be taken, then

we would be warranted in pursuing them. My goal in this section will be to sketch an

alternative answer to EVP based on the proposal in Chapter 1 that the epistemological

critical domain is properly framed by the constitutive goals of inquiry. This approach

differs from the bottom-up approach in a number of principled ways. I hope to make

these ways clear and suggest that the goals of inquiry give us a truly attractive option.

Ultimately, the explanatory power of the inquiry-based critical domain is to be seen as

some abductive support in favor of the view.

One might think the following: Look, it is just a primitive fact about knowledge that it is epistemically better than mere true belief. It is a fact that is *constitutive* of the epistemic evaluative domain and there is just no need to say anything more about it. If someone, for example, were to ask for an explanation of why delicious and aromatic coffee is better than bland and unsavory coffee from within "the world of coffee," the proper response might be nothing more than an incredulous stare. Why should we expect anything different in comparing the epistemic value of knowledge and mere true belief? There is, I think, some merit to this approach. It might, for example, help explain why we think that handcuffing knowledge's superior epistemic value to a particular metaphysical analysis of knowledge is to misconstrue its status within the domain. However, there is just no consensus to be found on what the epistemological critical domain is even supposed to be about. Many have the idea that it is supposed to be about some kind of success in cognition but, given that there are both monists and pluralists about fundamental epistemic value, veritists and non-veritists, we are left with some explaining to do. Given the disagreement, it would be best to present some plausible theoretical underpinnings for the value facts at hand.

In Chapter 1, I suggested that it is the constitutive goals of inquiry that act as the fundamental values within the epistemological critical domain. This view can now be put to work in answering EVP. In that chapter, I proposed the idea that there is a kind of inquiry called "propositional inquiry" and that it takes *knowledge* as its constitutive, ideal end. This led to VKT: For every proposition *p*, knowing whether *p* has fundamental epistemic value. We can now begin to see a relevant explanation for why tokens of knowledge are necessarily more epistemically valuable than tokens of mere true belief.

The nature of propositional inquiry dictates that it is knowledge that is the ideal end and that nature grounds knowledge's fundamental epistemic value. Mere true belief has a lesser standing (if it has any standing at all). Recall the case of Anne who propositionally inquires about whether *p* and luckily comes to have a true belief about whether *p* as a result of a knock on the head. (pgs. 27-9) I argued that Anne's belief, though true, is less than optimal vis-à-vis the goals of propositional inquiry. Her belief is *lucky* in a way that detracts from her success as an inquirer. She has not *tracked* the truth in the way that she would have if she came to *know* whether *p*. In that way, knowledge is better than mere true belief from the epistemic perspective. Further, knowledge has a *unique* epistemic standing being the optimal state, better than any other cognitive state that might fall short of it. For example, mere justified true belief, though perhaps an approximation of what we want when we propositionally inquire, it is still not the *ideal* end of inquiry like full-blown knowledge is.

Notice, now, that the answer to EVP given here does not appeal to any metaphysical analysis of knowledge. In that way, it is quite unlike the bottom-up solutions encountered earlier. Instead, the appeal is to facts about the goals of inquiry and knowledge's extrinsic (but necessary) relation to those goals. We need not have a particular metaphysical account of knowledge on the table and we need not reductively specify the *essence* of knowledge in order to see those relations. This is something that I think many theorists about epistemic evaluation have missed. If epistemic evaluation comprises a critical domain, then it need not be some feature of an analysis of knowledge that explains why it is necessarily better than mere true belief. More generally, when we

are trying to explain the relative domain-internal value of two different kinds of things, we do not always need to appeal to some real definition of those things.

Take an example: Suppose that I wanted to compare the domain-internal value of two different kinds of racket tensions from the perspective of competitive tennis. I want to know, say, why a racket strung at 55 pounds is better than a racket strung at 20 pounds. We would be travelling down the wrong road if we were to begin explaining such facts by comparing accounts of *what it is* for a racket to be strung at each tension level. In giving such an explanation, real definitions of particular tensile strengths would only play a role insofar as they help us specify each strength's *relation* to the relevant goals of the domain. What would explain the superior value of a racket strung at 55 pounds is an explanation of how certain physical facts about tensile strengths *relate* to the goal of winning tennis matches. A racket strung at 55 pounds allows for the right balance between power and control, better equipping a player to make winning shots. The relevant relations to the domain-internal goal play the most crucial role in such explanations and we can give them without completely specifying *what it is* to be a tensile strength of a certain magnitude.

Similarly, we can explain knowledge's epistemic superiority to mere true belief through the associated goals of the domain. The relative relations between each cognitive state and the goals of propositional inquiry seem to provide exactly the right touch. Knowledge is necessarily better than mere true belief with respect to those goals. As I mentioned in response to Marian David (see Chapter 1, pgs. 43-5), we should not conflate value theoretic considerations and metaphysical ones. This is especially easy to do when we are comparing the relative epistemic merits of knowledge and mere true

belief because the latter is typically taken to be a necessary condition on the former. This causes an automatic association with metaphysical analysis. Yet, insofar as epistemic evaluation comprises a critical domain, the value theoretic facts can occasionally be divorced from the metaphysics involved with the items of value. In this case, knowledge's extrinsic relation to the goals of inquiry gives us all we need in responding to EVP.

Supposing that this story about the fundamental epistemic value of knowledge is correct, what then about mere true belief (i.e., true belief that does not exhibit any further epistemic merit)? Does it always, like knowledge, have epistemic value? If the intuitions about Jimbo's beliefs are right, then it seems as though it does not. Can an inquiry-based critical domain meet the proper demands here? Bottom-up approaches begin with the fundamental epistemic value of true belief and then build up the value of knowledge through necessary derivative value relations. We are currently asking about what happens when we flip the script. We are asking about the status of mere true belief relative to the fundamental epistemic value of knowledge (perhaps best thought of as a "top-down approach" to the epistemic value of true belief). What kinds of necessary value relations might obtain between the two when *knowledge* plays the more central role?

A natural starting point in answering this question would be to survey some plausible necessary relations between mere true belief and knowledge and to see whether any of those relations generate legitimate derivative value. The relations already encountered in this chapter do not obviously obtain in any necessary way when we take knowledge to be the more fundamental value. Mere true beliefs at least initially do not

seem to have any necessary *instrumental* relation toward obtaining knowledge. After all, if a belief that *p* is a *merely* true belief, then that *rules out* the possibility that the belief that *p* instantiates knowledge. More, beliefs that are merely *based on* other true beliefs also will not instantiate knowledge, at least insofar as those other true beliefs lack some further epistemic merit. And even though mere true beliefs might sometimes be a *causal means* to the obtainment of knowledge in a way that is deviant vis-à-vis basing, such causal relationships only hold contingently, still allowing for cases of mere true belief that are not instrumental in this respect.[78] Mere true beliefs also do not seem to *manifest a respect* for knowledge or *manifest an achievement* of knowledge, especially in the way that Sylvan and Greco think of those relations. If there are necessary derivative value relations obtaining between knowledge and mere true belief, we will have to look elsewhere.

One might think that true belief, by itself, is epistemically good because it is an essential or constitutive part of knowledge and essential parts of good things are also good. Consider an example from Wayne Riggs: "If an MRI machine is valuable because

---

[78] There is a more general worry looming here about instrumental value production. Being a causal means to achieving the goals of inquiry might be thought of as too permissive with respect to the generation of epistemic value. For example, if gaining $1,000,000 is a causal means to doing more research and more comprehensively satisfying the goals of inquiry, then it should be considered as having some epistemic value. This, however, might seem strange. One might think that, whatever value the obtainment of $1,000,000 has, it would not be properly considered as having *epistemic* value. It should be noted, though, that whatever epistemic value obtaining $1,000,000 has under the current view, it would only be a highly contingent kind of value. There are many circumstances for which the obtainment of money would not be a causal means to the satisfaction of the goals of inquiry and, in those cases, it would not have any epistemic value at all. One reason why it might seem strange to think that it occasionally does have epistemic value is because epistemologists have been most interested in the *necessary* relations to epistemic goals and not contingent relations like being a mere causal means. When epistemologists *do* attend to causal relationships (e.g., reliable production, proper basing, etc.), more modal relations to the purported epistemic goals are posited. It would make sense for philosophers to focus in on necessary relations. After all, whether or not some particular state of affairs is a causal means to some other particular state of affairs is just an empirical question. I do not think, though, that this focus on necessary relations should make us worried about the more contingent varieties of derivative epistemic value. I will have more to say about instrumental value in the next chapter.

it helps save lives, then the individual components of it that make it work are also valuable." (2008, pgs. 316) Riggs calls this kind of mereological, derivative value relation "contributory" value. The issue, though, once again is that true belief is not literally a part of knowledge. Knowledge is not a mereological sum of true belief and other things and true belief does not "contribute" to knowledge in the same way that, say, an MRI machine's magnetic coils contribute to it. Insofar as being an MRI machine is good, so is being a certain kind of magnetic coil. We could physically replace a magnetic coil in an MRI machine with one of the same type and it would contribute toward the machine's operation in the same way. But, to speak in this way about mere true belief in its relation to knowledge would only be to engage in some lofty metaphor.

What could be more appropriately said about true belief in its relation to knowledge is that it is a *necessary condition* on knowledge.[79] However, is being a necessary condition on some good thing, by itself, always a relation that generates derivative value? Consider a moral example: Benevolent actions are morally valuable and being an action is a necessary condition on being a benevolent action. Does that mean that there is some derivative moral value accruing to any action at all in virtue of it being an action? Of course not. It seems as though, again, if there is some necessary derivative value relation obtaining between knowledge and mere true belief, we will have to look elsewhere.

An idea from Tamar Szabó Gendler (2002) presents a more promising strategy. When we evaluate certain kinds of entities, we occasionally treat "impure instances" of good things as relevantly similar to "pure" ones:

---

[79] Of course, this would only be true barring the worries brought out in footnote 49 (pg. 61).

> Suppose we venerate regular geometrical figures for their beauty, but certain approximations to regular figures also produce the same respect by way of resembling the ideal… [W]henever something is square-like, it is an appropriate target of geometrical veneration, and whenever something is not square-like, it is not an appropriate target of geometrical veneration; whether or not it is actually square has no bearing on its suitability as an object of geometrical veneration. But from this we are not entitled to conclude that it is square-likeness rather than squareness that explains the appropriateness of geometrical veneration… (pg. 47)

Supposing that squares are fundamentally good as a certain kind of ideal, certain approximations of squares might also be derivatively good in virtue of their square-likeness. Perhaps, then, supposing that knowledge is fundamentally good as an ideal end of propositional inquiry, mere true belief might also be derivatively good in virtue of its knowledge-likeness.

However, we need to be at least a little careful with this. If being a member of kind K is good, then there will be gradations of goodness that can be generated by K-likeness. Things can be *more or less similar* to other things in ways that are relevant to derivative value generation. For example, if squares are good in the way of geometrical beauty, then more derivative value will accrue to a nearly-squared rectangle than to a four-sided figure that is not nearly square at all. Of course, it is difficult to differentiate degrees of similarity in a principled way (e.g., see Deutsch 1998) but, in many cases where a certain gradational standard is fixed, it is relatively obvious that certain things are *greatly similar* to members of kind K, while others are *intermediately similar* to members of kind K, and while others yet are *minimally similar* to members of kind K, etc. Presumably, when derivative value is generated via K-likeness, the amount of value generated corresponds to the relevant degree of similarity. Only a minimal amount of value would accrue to entities that are only minimally similar to members of kind K, more would accrue to entities that are intermediately similar, and so on.

The relevant question, then, is the following: How similar to knowledge is mere true belief? Take, for example, Jimbo's true belief that his car is being towed. There are certain respects in which Jimbo's belief is similar to a relevant token of knowledge; they are both cognitive states, they both represent the world as being a certain way, they are both veridical, etc. Yet, there are many respects in which they are dissimilar. *Knowing* that one's car is being towed would have a particular kind of causal production, it would have a certain kind of justificatory status and relation to an agent's cognitive abilities, it would only endure in a circumscribed set of environments, etc.; all of these are features that Jimbo's mere true belief lacks. I would be tempted to say that mere true beliefs, those like Jimbo's that do not exhibit any further epistemic merit, are more like irregular trapezoids than they are like nearly-square rectangles with respect to square-likeness. They are only minimally similar to the ideal of knowledge. This would mean that mere true beliefs have only a minimal derivative value in virtue of their knowledge-likeness.[80]

This, however, seems to deliver exactly the right verdict, one that bottom-up approaches seem unable to replicate. From an epistemic evaluative perspective, how much value does Jimbo's belief have? As I have suggested, it has only a minimal amount, if it has any at all. It is hard to see how a bottom-up theorist could agree with this. It is somewhat telling that even the most dedicated bottom-uppers, despite their insistence on the fundamentality of true belief, often use the rhetoric of knowledge in describing the goals relevant to epistemology. Take, for example, Goldman (1999) on

---

[80] This is at least one major difference between the view espoused in these chapters and that suggested in Ahlstrom-Vij and Grimm (2013), where an "accuracy monism" is proposed. Though Ahlstrom-Vij and Grimm do suggest that epistemic value is framed by the goals of inquiry (rightly, I think), the fundamental values they propose still include mere true belief. Again, though, it is unclear as to how much of their own view can be mapped onto my own given that they do not make a distinction between genuine value and domain-internal value.

inquiry: "The primary purpose of asking a question is to learn the answer, the true answer, from the respondent. In asking someone for the time, or the location of the nearest post office, a questioner evinces a desire to *know* [my emphasis] something that she does not already know." (pg. 3) However, in his formal specification of the veritist position, he begins with a focus on knowledge and then he slips to mere true belief: "Veritistic epistemology (whether individual or social) is concerned with the production of *knowledge* [my emphasis], where knowledge is understood in the 'weak' sense of *true belief* [my emphasis]." (pg. 5) My suggestion is that we just own up to the fact that it is knowledge that is fundamental as an epistemic value and not true belief.

Given the considerations above, I take it that this more subsidiary role played by mere true belief is at least some abductive support in favor of the view outlined in Chapter 1.[81] Not only does that inquiry-centered epistemic evaluative practice provide an abstract, general solution to the most legitimate modal formulation of the value problem (i.e., EVP), it can render fine-grained verdicts about particular cases in a more satisfactory way than the most popular alternatives. There are, of course, many other epistemic desiderata that need to be handled by the theory (e.g., epistemic justification will be discussed in the next chapter) but, in accounting for the epistemic value-theoretic

---

[81] A different feature of belief that would potentially be difficult for the bottom-upper to handle is what I will call "mere safety." Take the following crude notion of safety: a belief that *p* is safe if and only if there is no close world surrounding the actual world where *p* is false. Suppose that Jimbo's belief that his car is being towed is not only true, but is also safe. (Perhaps, unbeknownst to Jimbo, there is a benevolent demon who always makes sure that Jimbo's beliefs about his car turn out to be true.) Jimbo's belief is *merely* safe in the sense that it is both true and safe without having any further epistemic merit. Now, if we were asked about the epistemic value of Jimbo's belief in this situation, it would be tempting to say that the mere safety involved does not contribute much, if any, epistemic value to the belief. The belief seems just as bad as his merely true belief discussed earlier. Yet, the bottom-upper might be committed to a different verdict. If true beliefs are of fundamental value, then it would seem as though safe beliefs (even merely safe beliefs) are epistemically better. After all, there is a sense in which the truth of safe beliefs is not easily lost and that would be a good thing given certain bottom-up commitments. If that is the case, I take it, again, that an inquiry-centered critical domain would fare better in these regards.

relations between mere true belief and knowledge, it sets a relatively high bar that any

alternative theory must meet.

Chapter 3: The Place of Epistemic Justification

This chapter is devoted to specifying the role of epistemic justification within the

critical domain outlined above. Since the domain is framed by the constitutive goals of

inquiry and since the goal of propositional inquiry is knowledge, the approach taken here

might appropriately be thought of as a "knowledge-first" account of the value of

epistemic justification. The first section of the chapter sets the stage for the positive view

to follow, outlining a number of theoretical desiderata while also specifying a way for the

view to remain neutral on some of the perennial debates about epistemic justification.

The second section outlines the positive view, arguing that epistemic justification is

derivatively valuable, exhibiting a kind of inquiry-relevant virtue. In forming a justified

belief, an agent "does her part" in gaining knowledge, contributing toward the goal by

exercising a cognitive capacity. The third section further specifies the view while

fielding a number of worries that one might have about the proposal. These include

worries related to doxastic involuntarism, epistemic demands involving evidence

gathering, the proper incorporation of negative valuation, and Selim Berker's arguments

against teleological approaches to epistemic justification.

3.1: Knowledge and Epistemic Justification

In the previous chapter, it was argued that true belief does not have the important

standing that it is often thought to have within epistemology. Knowledge, not mere true

belief, has fundamental epistemic value. This is not to say that there is nothing to the

slogan that we "aim at the truth" in our distinctively epistemic endeavors; it is merely to

say that we cannot interpret this slogan as meaning that we simply aim at true belief. We look to grasp the truth in a stronger sense than merely believing it. The lesson is that, if our comparative evaluations of epistemic worth are to make any sense, we need to own up to the fact that, by focusing on true belief, many epistemologists have misidentified what is really at the heart of such evaluations. Our attention needs to be redirected.

This is an important consideration as we approach the topic of epistemic justification. Traditionally, being justified has been thought to be an epistemically good-making feature of belief, even if its goodness is derivative to the goals involved in "getting at the truth" in whatever sense. The claim that epistemically justified beliefs are epistemically good ones does have a vacuously true ring to it. The language of "justification," by its very nature, is at least evaluatively loaded if not outright deontic. Perhaps, though, we should not be so quick about the normative legitimacy of our justification attributions. After all, if we can be wrong about the value-theoretic status of true belief, then perhaps we can also be wrong about the status of the different features that have been seen to constitute epistemic justification, especially if those features are only significant in virtue of their relation to true belief.

At least intuitively, however, there are some important differences in the way we see the value of truth and the value of justification from the epistemic perspective. Take, for instance, the fact that we do not seem to place much epistemic worth in *merely* true beliefs. Recall Jimbo who luckily comes to a true belief about his car being towed on the basis of a tarot card reading. Such a belief does not seem much better off, epistemically speaking, than a false one formed in the same way. The same does not hold when we compare *merely* justified beliefs (i.e., those beliefs that are justified but do not exhibit any

other kind of epistemic merit) with their unjustified counterparts. Suppose that Jimbo

forms the belief that his car is being towed on the basis of some justifiable grounds. Even

if his belief turns out to be false and thereby does not constitute knowledge, his belief is

still epistemically better than it otherwise would be if it were not so formed.[82]

Given the approach taken in these chapters, if the intuition that beliefs have

epistemic merit in virtue of being justified is to be vindicated, then it will have to be

explained how justification plays some important role in our evaluations structured by the

goals of inquiry. Moreover, given the previous chapter, a discussion about justification's

connection to true belief would not be enough. Perhaps a truth-connection will end up

playing some part in a story about the role of justification, but it could only do so if there

is a further narrative, one that can adequately tie such a truth-connection up with the most

fundamental goals involved in inquiry. Since it has been argued that the goal of

propositional inquiry is knowledge, the undertaking in this chapter might be seen as a

"knowledge-first" approach to the epistemic value of justification, not in the sense that

knowledge is somehow metaphysically prior to justification but that the salient epistemic

value-theoretic facts about justification will be explained in terms of the value of

knowledge (i.e., the epistemic value of justification is derivative to that of knowledge).[83]

One issue that should be clarified at the outset is the following: Ordinarily, when

we talk about the value of epistemic justification, we take it to be a good-making feature

of *beliefs*. This means that the real explanandum is the value of *doxastic* justification and

---

[82] This claim assumes that it is possible for a token belief to be simultaneously epistemically justified and false. I will be assuming this to be possible throughout the chapter. See Littlejohn (2012) and Steglich-Petersen (2013) for contrary positions, though I will not have the space to discuss them in detail here.

[83] Contrast this with the view developed in Kelp (forthcoming).

not merely *propositional* justification.[84] Depending on one's theoretical purposes,

propositional justification and its relation to other forms of justification or knowledge

may be worthy of an extended discussion and may play some important role in

explanation. Here, however, our focus will be on the evaluative status of beliefs that

have actually been formed and are actually maintained in agents and on the ways that

such beliefs have been formed or maintained, not on the evaluative status of a state that

an agent is in when they *would* be justified in believing a proposition if they were to

properly capitalize on its support. While it may be epistemically valuable to be

propositionally justified with respect to some proposition *p*, I take it that such value is

derivative to *actually* capitalizing on that support in coming to be doxastically justified in

believing that *p*.

With that being said, there is a dizzying array of theories available to us about

what doxastic justification consists in. Some theorists are access internalists, claiming

that the justificatory status of a belief wholly depends on whether a believer has properly

based her belief on mental states that she has some kind of special, reflective access to.

Others deny such claims and are thereby access externalists. Some theorists claim that

justificatory support must be foundational and directional in structure while others claim

that it is more like an interconnected web of symmetrical support relations. Some

theorists put forward positive theories, requiring that justified beliefs be based on good

reasons. Others put forward negative theories, requiring of a belief only that certain

defeating conditions be absent. Some theorists claim that justification consists in

---

[84] See, e.g., Firth (1978) and Turri (2010) for more on the distinction between propositional and doxastic justification. The considerations below will also apply to other, related distinctions like that of *ex ante* and *ex post* justification (see Goldman 1979 and Ichikawa and Jarvis 2013, pg. 163).

something deontological while others require it to play some "epistemizing" role in closing the gap between mere true belief and knowledge. The array of theories available to us is so dizzying that some philosophers have even wondered whether the alternatives are even aimed at the same object, whether they are even competing theories about the very same phenomenon. According to William Alston (2005), we are best off dispensing altogether with the term "justification" and focusing more directly on the various epistemic desiderata that often undergird our applications of it.

I cannot even hope to settle many of these disputes here. In fact, I will be trying to remain neutral on as many of these issues as I can, including Alston's contention that the traditional debates about justification are not even getting on about the same thing. What I can do, however, is discuss a theoretically useful way of thinking about epistemic justification and discuss the kind of role that justification might play in the epistemic evaluative practice theorized about above. My hope is that, at the very least, I will be circumscribing an important property of beliefs in a way that will redeem justification as occupying a valuable role vis-à-vis inquiry. My general approach follows Declan Smithies' recommendation (2015, pgs. 224-5) that, in theorizing about justification, we should antecedently pick out a purpose for using that concept in our epistemic evaluative practice before we ask what something must be like to play that role or what playing that role must consist in. Put differently, I will be outlining a baseline, guiding principle for us in our application of the concept of justification, one that will capture its axiological status within the framework provided, and I will do so prior to identifying the specific nature of the thing that falls under that concept. In that way, I hope to explain the

epistemic value of justification while at the same time, for example, skirting the debate between internalists and externalists. I will have more to say about this below.

I take it as a given that the value of justification in the epistemic evaluative domain is derivative, in at least some way, to the goal of "getting at the truth." If our interests are epistemic, then it is a good thing to have epistemically justified beliefs. Yet, we do not seem to desire such things fundamentally and for their own sake. A *merely* justified belief, while epistemically good, leaves something to be desired and the traditional way of explaining this is by appealing to its relation to some further epistemic good. I will not be departing from this tradition. For many, however, true belief plays the role of this further good. Here, the value of justification must alternatively be explained in its relation to the constitutive goals of inquiry (more specifically, propositional inquiry). Moreover, justified beliefs seem to have epistemic value even when that further goal is not achieved, even when knowledge does not obtain. The challenge will then be to specify some derivative value relation that can do justice to these ways of thinking about justification.

To make the explanatory desiderata clearer and to motivate a connection that I will eventually make, consider an example (call it "Base Case"): Anne is arriving at home from work and that she is wondering whether her husband is also at home. She pulls into her usual parking spot in the driveway and sees that her husband's vehicle is parked in its usual place on the other side. As she walks inside the house, she notices that her husband's keys, his wallet, and his phone are all laying on the kitchen counter, the usual place that he deposits his belongings when he arrives at home from work. She notices that the television is on in the room adjacent to the kitchen and that one of her

husband's favorite shows is on. Though she cannot directly see any particular person, she does hear someone munching on a snack in the armchair in front of the television. On the basis of all of this information, Anne forms the belief that her husband is, indeed, at home.

Supposing that there are not any other defeating justificatory considerations present, we would typically take Anne's belief to be epistemically justified, at least to an important degree, and the fact that she has a justified belief is an epistemically good thing. Her belief has some epistemic merit. The interesting thing is that her belief has that valuable status independently of some further conditions that may or may not obtain. For instance, her belief has some epistemic value even if it turns out to be false. Perhaps all of the evidential considerations specified above are highly misleading and, because of that, Anne has come to falsely believe that her husband is at home. Given the evidence, it is at least possible that Anne's husband walked a block down the street to talk to some neighbors before she arrived home and that someone else is responsible for the munching noises in front of the television, perhaps a family friend with a house key. The point is that, even if this is the case, the justificatory status of her belief is still an epistemically good-making feature of it.

Considering other more bizarre logical possibilities, even if Anne's belief turns out to be true, it might be Gettierized in a way that would preclude it from constituting any sort of knowledge. Perhaps, again unbeknownst to her, she was kidnapped overnight by some evil, Gettier-case fetishists who are feeding her a simulation of experiences as of the situation described above. As things turn out, her husband may still actually be at home outside of the simulation making her belief true, even though that belief would

seemingly not constitute knowledge. Even if this is the case, Anne would still be justified in believing that her husband is home and that would still be a good-making feature of her belief from the epistemic perspective. Justification seems to have epistemic value even if the further goal involved in "getting at the truth" does not obtain. Surely it would be epistemically *preferable* for her justified belief to also turn out to be true and for it to constitute knowledge, but it still has some worth even if this does not turn out to be the case.

Hence, even *merely* justified beliefs have a significant sort of epistemic value. Since the same thing could not be said about merely true belief, we seem to have a phenomenon that calls for an explanation. Given the sort of knowledge-first approach that will be taken here, a difficulty arises in the following sort of way: We have already seen that being a necessary condition on something of fundamental value does not automatically generate any sort of derivative value (see Ch. 2, pgs. 98-9). Even if it did, however, and even if justification is a necessary condition on knowledge as the traditional JTB+ framework would have it, being a "part" of knowledge in this way cannot do all of the relevant value-theoretic explanatory work since the very same thing can be said about the truth condition. We need to set the justification condition apart from the truth condition and not merely appeal to their similarities. This also means that we should avoid the strategy used in the previous chapter to capture the derivative epistemic value relations between mere true belief and knowledge (pgs. 99-101). After all, a merely justified belief hardly seems to approximate the ideal kind of knowledge in inquiry any better than mere true belief. The primary focus of this chapter, then, will be on attempting to pick out a special feature of justification in its relation to knowledge that

would be able to explain its distinctive sort of standing within the epistemic evaluative practice.

Before specifying a positive account, it is worth briefly discussing another strategy that I do not find to be particularly promising. In the last chapter, a worry was brought out about *instrumental value relations* and their potential role in accounting for the superior value of knowledge over and above mere true belief (see pgs. 64-6). As Linda Zagzebski and others argue, if epistemic justification is only derivatively valuable in a way that it instrumentally promotes true belief, then it cannot account for the extra value that seems to accrue to a belief in virtue of its being justified. Using a reliable coffee machine as an example, Zagzebski points out that if a good cup of coffee happens to be produced by this reliable machine, that fact does not make the good cup of coffee any better. The relevant value of the product cup of coffee is exhausted by the value of its taste and aroma properties, those properties that the reliable machine is instrumental toward. Moreover, the fact that a *bad* cup of coffee happens to be produced by an instrumentally valuable machine does not make that bad cup of coffee any better either. Similarly, if epistemic justification is only instrumentally valuable vis-à-vis true belief, then the fact that a belief was produced in a justifiable way would seemingly add no value to it, no matter whether the belief is true or false.

Even though Zagzebski's example is aimed at reliabilist theories of epistemic justification in particular, I argued that a similar point can be made about *any* theory of epistemic value that takes justification to be merely instrumentally valuable relative to true belief, including more internalist theories (see pgs. 71-2). Here, I would like to argue that a similar sort of constraint also saddles knowledge-first approaches to the

value of epistemic justification. The issue for veritists like Goldman was that true beliefs do not get any better by being justified (and neither do false beliefs) if justification is merely instrumentally valuable. More generally, if some belief is an instance of fundamental epistemic success (or failure), no instrumental relations involved in the production of that belief could make it, as a product, any better off.

Since we are now considering the hypothesis in which *knowledge* has fundamental epistemic value, it will only be cases of a justified belief's *failure* to constitute fundamental value that will be relevant. It makes little sense to demand an explanation for why some bit of knowledge (the relevant fundamental epistemic success) might be better off when it is justified as opposed to when it is not (especially if justification is taken to be a necessary condition on knowledge).[85]   However, as we have seen at the outset of this chapter, we *can* demand an explanation for why beliefs that fail to constitute knowledge are epistemically better off when they are justified. This strikes us as a genuine phenomenon. In other words, beliefs that fail to constitute a fundamental epistemic success get epistemically better in virtue of their being produced justifiably and we need to somehow explain that fact. Now, suppose that we attempted to explain the value of justification in terms of some kind of instrumental relation toward knowledge. Zagzebski's problem would arise once again. If we think of justification in that way, it is not clear how a product belief that fails to constitute knowledge could be any better off by being produced in a way that is merely instrumental to generating knowledge. In

---

[85] For some exceptions to the orthodoxy that justification is necessary for knowledge, see Goldman (1967), Armstrong (1973), Dretske (1981), and Kornblith (2008). Here, I will be taking justification to be a necessary condition on the sort of knowledge that is of fundamental epistemic value. Philosophers who deny that justification is necessary typically wish to specify a sort of cognitive success that might be had in common among human adults, children, non-human animals, etc. and, hence, they have a lower standard in mind, not some ideal epistemic relation that might be had between an agent and a true proposition. See footnote 26 in Chapter 1.

other words, if epistemic justification is valuable merely as an instrumental means to knowledge, we will not have explained how a non-knowledge-constituting belief that is produced justifiably is any better than a belief that is produced in an unjustified way. We will not have explained what seems to be an important fact about epistemic value.

For this reason, I will leave purely instrumental approaches to justification behind. An instrumentalist approach that would be tempting to carry out would be one in which justification is *probabilistically* related to knowledge and, hence, instrumental toward it, perhaps as a way of ruling out certain kinds of luck or by having some probabilistic relation to truth. Many philosophers have thought of epistemic justification in these ways, as playing some sort of probabilistic and instrumental role in our cognitive economies (see, e.g., Goldman 1979 and Swain 1981).[86] The point here is *not* to deny that epistemic justification has some probabilistic relation to truth or knowledge, it is merely to say that, in order to fully account for it as the distinctive kind of status within the epistemic evaluative practice that we typically take it to have, some further explanation is needed. Probabilistic and instrumental relations, by themselves, will not do as adequate derivative value relations.

3.2: Doing One's Part in Inquiry

Consider the following case (call it "The Shot"): During a particular point in the finals of the Wimbledon championships, Serena finds herself in what she thinks is a position to hit a winning shot against her opponent who is approaching the net. Her footwork, her timing, and her execution in general are impeccable as she attempts a

---

[86] Of course, such probabilistic connections are not without their detractors. See, e.g., Foley (1987, pg. 156) and Chisholm (1989, pg. 76).

passing shot. The ball sails down the line with both pace and accuracy and it has a heavy topspin that would cause the ball to kick away from her opponent if it were to strike the ground. Unluckily, though, Serena's opponent happens to correctly guess the exact destination for her shot down the line. As Serena is striking the ball, her opponent fortuitously scurries to the side of the court that allows her to (barely) get her racket on the ball. The volley then clips the top of the net and dribbles over onto Serena's side, ending the point. Though disappointed, Serena readies herself for the next point knowing that her shot was a good one. Her technique was sound and there was no way for her opponent to have known where the ball was going to go. Her shot was an exhibition of her capacities as a competitive tennis player and she fully did her part in hitting a winning shot, even though the unfortunate circumstances of the situation caused the goal of the shot to be thwarted. She wishes she could have done more to achieve the goal of the shot but, alas, she could not have. The failure was out of her hands.

Consider a second case (call it "The Swim"): Robert, a lifeguard, sits atop his tower post at Ocean Beach and recognizes that an apparently unexperienced swimmer is being swept under the water and away from shore by a rip current. He reacts quickly, sprints toward the water, dives in, and swims as vigorously and efficiently as possible toward the drowning patron. The rip current, however, continues to drag the patron further and further away from shore, making it more and more difficult for Robert to reach him. After a significant amount of time, Robert reaches the unconscious patron, pulling him out of the tidal current and begins working his way back toward the shore. By the time he reaches the shore, however, the patron is not breathing and has no pulse. Exhaustedly administering CPR, Robert realizes that the patron is still unconscious and,

despite his best efforts, he just cannot seem to revive the drowned citizen. By the time

additional emergency assistance arrives, the patron has passed away and the time of death

is recorded. As Robert thinks back on this horrible situation, he wishes that he could

have done more for the patron in his life-saving attempt. Yet, he could not have. He did

his part as a good lifeguard and his attempt was an exhibition of his good lifeguarding

capacities. A lesser lifeguard would not have even been able to reach the patron, let

alone get him all of the way back to shore against the rip current. It was an excellent life-

saving attempt by Robert, even though the goal involved in his attempt was thwarted by

conditions that were out of his hands.

There is a really natural way of drawing some similarities between the situations

involved in The Shot and the The Swim. In both cases, an agent occupies some role or

position that has a certain telos or goal associated with it. Serena is a competitive tennis

player and, as such, she has a goal to hit winning tennis shots and, ultimately, to win

tennis matches.[87] Robert is a lifeguard and, as such, he has a goal to keep his patrons

safe from harm. Because they each occupy their own respective positions, it is

appropriate to evaluate them, their activities, and other relevant items with respect to

those associated goals. In other words, even though Serena and Robert's situations are

very different, critical domain evaluation is equally appropriate.

What is also relevant for the purposes here is the fact that Serena and Robert both

ultimately *fail* to achieve some relevant goal within their respective domains. Serena

fails to hit a winning shot and, similarly, Robert fails to save a patron. Yet, as the cases

go, there seems to be some sense in which it is correct to say from the perspective of the

---

[87] The goals involved in competitive tennis playing might not be so easily specifiable. See footnote 1 in
Chapter 1.

relevant domains that Serena's shot was still a good one and that Robert's life-saving attempt was good. It is this feature of The Shot and The Swim that I think is worth focusing on in an attempt to explain the special sort of domain-internal value that accrues to epistemic justification. What ultimately explains why Serena's shot is a good one are facts about how the shot exhibits her capacities as a good competitive tennis player in a way that she "does her part" in achieving a relevant goal. Similarly, what explains the goodness of Robert's life-saving attempt are facts about how the attempt exhibits his capacities as a good lifeguard in a way that he "does his part" in saving the patron. If epistemically justified belief can properly be thought of in a similar way, then perhaps we will have the right kind of derivative value relation to do the work needed.

Consider, once again, Base Case (pgs. 108-10). Anne is inquiring about whether her husband is at home and she comes to believe that he, in fact, is. Now, as an extension to Base Case, suppose that she is wrong about this and that the person in front of the television really does turn out to be a family friend with a key to the house. At least for the moment, Anne has failed to achieve the goal involved with her token propositional inquiry. Yet, given that her belief is justified and that she believes on the basis of the (highly misleading) considerations outlined above, her belief does seem to exhibit a relevant kind of capacity, a capacity that makes the way that she forms her belief good. Had she come to a different belief on the basis of the very same considerations that she was entertaining, she would appropriately be seen as a defective inquirer (and similarly if she suspended belief). She, for the moment, does her part in achieving the goal of inquiry but the achievement of that goal has been thwarted by conditions that are out of

her hands.[88] She contributes toward gaining knowledge by meeting a justification condition, and yet the world has not cooperated. Hence, in a similar way that Serena and Robert fail to achieve a relevant goal within their domains but yet, in doing so, they produce a good shot and a good life-saving attempt, respectively, I want to suggest that Anne's belief should also be seen as good, as having some kind of domain-internal value.

The more general story to be told is this: Within a critical domain that is framed by an activity or role with a particular telos, if an item of evaluation in that domain exhibits a relevant capacity associated with that role in a way that an agent does her part in achieving something of domain-internal value, then that item also has domain-internal value.[89] So, Serena's shot has domain-internal value because of the way it exhibits her capacities as a good competitive tennis player in contributing toward hitting a winning shot. Robert's life-saving attempt has domain-internal value because of the way it exhibits his capacities as a good lifeguard in contributing toward the safety of a patron. Anne's belief has domain-internal value because of the way it exhibits her capacities as a good inquirer in contributing toward gaining knowledge. Each agent exercises a good-making capacity in their respective activities and, in doing so, they do their part in achieving some domain-internal good. This makes the various items that are under evaluation (i.e., the shot, the life-saving attempt, and the belief) also good. Doing one's

---

[88] This would also be the case if her belief were justified and true, yet Gettierized because of certain unlucky environmental conditions. In such a case, again, the failure would properly said to be "out of her hands."

[89] It should be noted that not every critical domain will have this sort of structure, being framed by some teleological activity or role. All of the domains mentioned above are framed in such ways: competitive tennis playing, lifeguarding, and inquiring. Other domains, however, have no teleological activity to appeal to in specifying their most fundamental values. For example, the most fundamental values within Sosa's "world of coffee" seem to be best specified directly as certain taste and aroma properties of coffee, not via some activity or role that constitutively takes some goal to be achieved.

part in this way is a *derivative value relation*, derivative relative to the more general goals involved in the roles that frame each domain.

This approach to the axiological status of epistemic justification is most properly thought of as virtue epistemological. The source of the value of epistemically justified belief is the capacity exhibited by an agent in meeting certain standards of excellence; in this case, excellences associated with inquiry. Since the goal of propositional inquiry is knowledge, an agent can only do so much in contributing toward the satisfaction of that goal.[90] She can believe justifiably and, in doing so, she may end up gaining knowledge. However, whether she gains knowledge or not will depend on certain other factors independent of her own contribution. For instance, the basis of her belief might end up being highly misleading, supporting some false proposition, thwarting the achievement of the goal. And even if the basis is not misleading, a justified true belief might be Gettierized, also thwarting the achievement of knowledge.[91] In any case, the justified believer still excels by way of the standards of good inquiry and it is this kind of exhibited excellence that explains justification's epistemic value.

Given the *prima facie* plausibility of this kind of story, it might also be plausible to think that the very role of the concept of epistemic justification is to capture this kind of mental contribution from an agent toward gaining knowledge, abstracting away from

---

[90] Simon Blackburn (2001) suggests that we see the most fundamental epistemic goals as parallel to *eudaimonia* as the goal of living. Like an account of virtue ethics, we should then expect that there will be certain excellences and exhibitions of capacities by agents that will contribute toward those goals.

[91] Although it *is* being suggested that knowledge is something like JTB+, there is no implication that such an analysis would be a successfully reductive one. For all that has been said here, the condition required over and above JTB may not be specifiable in a way that does not refer back to knowledge itself. Because of this, the framework set out above can be made more or less friendly to accounts like that of Williamson (2000).

other facts that would need to obtain in order to be successful.[92] We can think of the

concept of justification through some baseline principle like this and yet remain neutral

on what, exactly, it is for some feature of a belief to play this role. Even though a belief's

being justified indicates that an agent has done her part vis-à-vis the goal of propositional

inquiry, what exactly this amounts to remains to be filled in. Hence, the derivative value

relation outlined above allows us to remain neutral on many of the perennial debates

about epistemic justification. An access internalist, for example, could take this relation

on board and still argue that what it is to satisfy it, to "do one's part" in this way, is for

one to properly base her belief on facts that are reflectively accessible to her. An

externalist will deny this, of course, but she is also free to think of our attributions of

epistemic justification as indicating whether or not an agent has contributed toward

knowledge in some necessary way.[93]

      Take, for instance, the popular view that doxastic justification requires an agent to

properly base her belief on certain kinds of supporting conditions. A plausible

interpretation of this requirement is that it involves an agent *doing something well*.

Proper basing is a kind of contribution on the part of the agent, an exhibition of a

---

[92] Take, for example, the way that Alexander Bird (2007) describes the role of justification: "Anything short of knowledge is failure. But some failures are worse than others. And in particular some failures can be laid at the door of the believer, because the source of failure is one or more of the believer's mental states, and some failures can be ascribed to mischance, in that the failure is due to some mentally extraneous factor. The role of the concept of justification is to mark the difference between these different sources of failure." (pg. 95) Also see Whitcomb (2014).

[93] Externalists who stress the importance of reliable sub-personal systems will not think that the contribution by the agent need be executed *at the level of the agent*. One might argue that the contribution toward knowledge is often exhibited below the level of conscious awareness. Nonetheless, the analogies given still hold. This kind of story would be perfectly consistent, for example, with a story about Serena's good shot in which her sub-personal systems were responsible for her contribution toward success. Tennis players do not exercise agency in every movement they make on the tennis court and, yet, the sub-personal systems that are responsible for those movements might exhibit tennis playing skill or not, appropriately contributing toward the goal or not. Even internalists who think that we can form justified beliefs below the level of conscious awareness can endorse an analogy like this.

capacity, and we can say this prior to taking a position on whether the supporting

conditions that the beliefs are based on must be "internal" to the agent or not. We see

accounts of doxastic justification like this across the internalist-externalist spectrum. For

example, Conee and Feldman's (1985) internalism utilizes the basing relation in a

working out of an evidentialist theory of epistemic justification, as does McCain's (2014)

more recent iteration. Similarly, Alston's (1988a) "internalist externalism" requires

justified belief to be "based on an adequate ground." (pg. 265) Further down the

externalist line, Goldman's (1979) process reliabilism utilizes the basing relation in

specifying the kinds of paradigmatic processes that confer a positive justificatory status.

Because of this focus on the basing relation, it could easily be argued that all of these

philosophers see epistemic justification as indicating a sort of contribution toward

knowledge on the part of the agent but, yet, disagree about what that contribution consists

in.[94]

Since the account above is an instance of a virtue epistemology, a note should be

made about how it differs from other virtue approaches that are somewhat similar. For

---

[94] Two relevant points should be made here: (1) One might worry that certain accounts of epistemic justification that do not focus on a basing relation would be less friendly to the baseline principle suggested above. For example, Harman (1986), Swinburne (2001), and McGrath (2007) all put for forward negative theories of justification that do not require positive instances of proper basing. Again, though, I see no reason to think that a dispute between positive and negative theorists is not downstream to conceiving of epistemic justification as an exhibition of a capacity on the part of the agent. Perhaps the relevant contribution on the part of the agent only requires her to ensure that no defeaters are present for her belief. In that case, what it would take for an agent to "do her part" would be to ensure such things and nothing more. (2) In saying that doxastically justified beliefs exhibit a kind of contribution toward knowledge, the possibility of *innate knowledge* is not to be thought of as automatically being ruled out. For example, philosophers like Peter Carruthers (1992) can agree with what has been said here and still think that we have a kind of innate knowledge of, say, the principles of folk-psychology. If we have such knowledge at birth, we (including our sub-personal systems) still might make a kind of contribution in its origination and maintenance. For all that has been said here, that contribution might just consist in the fact that the belief is caused by certain reliable cognitive processes generated by natural selection. More internalist-minded philosophers are unlikely to think that such beliefs are justified and are instances of knowledge, but that is precisely because they think that the contribution on the part of the agent must be more robust. Again, the current thesis is intended to stay neutral on a debate like this.

example, Sosa's (2007, Ch. 2) virtue epistemological account is comparable to the one here as he utilizes an "AAA framework" for belief. For him, justified beliefs are good because they are "adroit" beliefs (i.e., those that exhibit a skill on the part of the agent or her sub-personal systems in bringing about accuracy).[95] In the same way that an archer's shot can be accurate without being adroit and adroit without being accurate, a belief can be true without being justified and justified without being true. Moreover, in the same way that an archer's shot can be both accurate and adroit and yet fail to be "apt" (i.e., accurate *because* adroit), a belief can be both true and justified and yet fail to be an instance of knowledge. In fact, apt belief is what Sosa takes knowledge to be (belief that is accurate *because* it is adroit).

Though the account presented here agrees with Sosa's in a number of respects, there are a couple points of disagreement. First, I have described justified belief as the exhibition of a contributory *capacity* while Sosa frames it in terms of the exhibition of a relevant *skill*.[96] There is a reason for this choice in terminology and it is related to a worry brought out by Robert Audi (2001), one originally directed more generally at virtue accounts. Audi claims: "It would seem possible to have a justified belief when no associated trait deserving the name 'virtue' is present. A generally uninsightful person can suddenly see a pattern (in a way we think of as indicating insight) and thereby be justified in a belief about it." (pg. 89) I take Audi's point to be that forming a justified

---

[95] It is actually relatively rare for Sosa to speak overtly about his adroitness condition as being an account of justification. He does, however, make this connection explicit on certain infrequent occasions, indicating that he does have something like justification in mind (see, e.g., Sosa 2007, pg. 42).

[96] I also take this point to betray a preference for Greco's (2010) virtue reliabilist language over Sosa's as Greco talks about the relevant virtues in terms of "abilities." Generally speaking, I see no reason why "ability" and "capacity" could not be thought of as synonymous. This is not, of course, to say that the account given here is in full agreement with Greco; see, e.g., the extended discussion of Greco's views in Ch. 2 (pgs. 81-92).

belief is possible for pretty much anyone, even for those who are pretty generally unskillful in their cognitive lives. According to Audi, this should make us question whether any sort of virtue need be exhibited in the formation of a justified belief. Suppose that, in Base Case, Anne is generally not the kind of person who thinks carefully about her evidence and she usually does not form her beliefs on the basis of genuinely good epistemic reasons even though, in this instance, she happens to do so. Would it really be right to say that her belief that her husband is at home exhibits some kind of virtue, making it epistemically valuable?

Audi's worry here is perhaps best framed as one about the possibility of *lucky* justified beliefs. The account presented here seems to avoid this worry in a way that Sosa's does not. It allows for lucky justified beliefs. In the example, while Anne is not particularly *skillful* in the relevant way since she does not really seem to possess the relevant habits or dispositions, she still does exercise an inquiry-relevant *capacity* as she is still able to "do her part" in contributing toward knowledge in a way that a good inquirer would. Something very similar could be said if we replace Serena and Robert in The Shot and The Swim, respectively, with less skillful individuals within their domains. Suppose, for example, a novice tennis player performs the very same movements in the very same way that Serena does. Even if this novice is not skillful like Serena, she still exercises the very same contributory *capacity* that Serena does as she "does her part" in the way that a good tennis player would, at least insofar as she really does all the very same things in hitting the shot (i.e., hitting the ball with the very same pace, accuracy, and spin, checking the opponent's position on the court, disguising the direction of the shot, etc.). That is enough to make her shot a good and virtuous one from the perspective

of competitive tennis playing, even if the novice is not particularly skillful. Similarly for justified believers.

Another way in which the account presented here differs from Sosa's is that the relevant kind of virtue, for Sosa, is merely truth-aimed while in the account above it is knowledge-aimed. Sosa's AAA framework makes the virtue manifested in a belief derivatively valuable relative to the more fundamental goodness of accuracy (i.e., true belief). We can see this by utilizing the ontological dependence test for derivative value applied in the previous chapter (pgs. 63-4). The value of Sosa's adroitness asymmetrically depends on the value of accuracy, making it derivatively good relative to true belief. The previous chapter was devoted to showing how this strategy is not a particularly promising one and that we should instead focus on knowledge as the relevant fundamental epistemic value. It is this knowledge-first feature of the account above that makes it uniquely suited to meeting our explanatory demands. The kind of capacity exhibited in doing one's part to achieve the goal of propositional inquiry (i.e., gaining knowledge) puts the fundamental epistemic value in the right place with respect to those considerations.[97]

3.3: Some Worries and Replies

It would be quite natural to associate this language of "doing one's part" with certain other similar turns of phrase like "doing one's duty" or "meeting one's responsibilities." The first worry that I will consider in this section is related to these

---

[97] Lisa Miracchi (2015), a Sosa student, also suggests making the shift to a more knowledge-centric brand of virtue epistemology, though her motivations to do so are related to her Williamsonian leanings about knowledge being a mental state in its own right, independently of belief.

ways of describing the derivative value relation outlined above. There is a sense in which I think that we should be comfortable thinking of an agent's "doing one's part" as an agent meeting what might be called her "role responsibilities."

Here, I use the term "role responsibility" in a stipulative sense that may only loosely connect to more ordinary uses of the phrase. By "role," I only mean to specify the general kind of position that an agent occupies when engaging in an activity with a telos. Competitive tennis players occupy roles in this sense, lifeguards occupy roles, inquirers occupy roles, etc. In evaluating agents with respect to such roles, there are certain kinds of things that the agent can, in a sense, properly said to be responsible for in contributing toward the associated end.[98] Achieving the goal of an activity often requires a contribution from an agent along with the obtaining of certain independently fortuitous states of affairs. When an agent *does* contribute in the relevant ways, she meets her role responsibilities regardless of whether or not those independently fortuitous states of affairs obtain and regardless of whether or not the end associated with the role is ultimately achieved. In this stipulative sense, Serena meets a kind of role responsibility with respect to competitive tennis playing, even though the goal of her shot was not achieved. Robert meets a kind of role responsibility with respect to lifeguarding in his life-saving attempt. Inquirers meet a kind of role responsibility by forming justified beliefs, independently of whether or not those beliefs turn out to constitute knowledge.

---

[98] In this way, the notion of role responsibilities employed here is wider than that used by H.L.A. Hart (2008, pgs. 212-4) who thinks of such responsibilities as attaching only to the duties associated with positions in existing social organizations and institutions. Here, so long as an agent may engage in a particular activity with a telos and contribute toward the achievement of that goal in a certain way, then there will be a way in which that agent may or may not do one's part with respect to that role independently of whether there be any social scaffolding beneath it. Hart's role responsibilities are a proper subset of the kind of role responsibilities discussed in this context.

The point of the previous section was that items that exhibit an agent's meeting her role responsibilities in this sense are good vis-à-vis the critical domains framed by those roles.

Yet, we need to be careful with such language. In speaking about epistemic justification as "doing one's part" or "meeting one's role responsibilities," an immediate worry springs to mind given that the central topic under discussion is the formation and maintenance of belief. The language involved here might easily be associated with deontological concepts like that of blameworthiness, praiseworthiness, and even retributive notions like desert. "Responsibility," in particular, often has such connotations, especially in the context of discussing *moral* evaluation.[99] According to many epistemologists, however, such concepts cannot be applicable to our doxastic practices. This is because, unlike actions, we purportedly do not have the sort of control over our beliefs needed to merit such application. Many have thought that, if even a modest form of doxastic involuntarism is true, then epistemic justification cannot be a deontological kind of evaluation.[100] Given the relative plausibility of involuntarism (for at least many of our beliefs), perhaps an account of the value of justification in terms of "doing one's part" or in terms of meeting one's inquiry-specific role responsibilities is bound for failure.

In responding to this worry, it would be ideal to stay neutral in the debate about doxastic voluntarism (at least here, though I will discuss some relevant positions within the voluntarist debate in the next chapter). Hence, it will be important to more clearly identify what role responsibilities are supposed to be in this context and how they might

---

[99] See Slote (1990), however, for an excellent discussion about how certain kinds of moral evaluation are independent of this kind of responsibility.

[100] See Alston (1988b) and Feldman (2000; 2001) for extended discussions on this issue.

apply independently of considerations about freedom and conscious control. Take the case described in The Swim: Robert does his part in his life-saving attempt and meets his role responsibilities vis-à-vis lifeguarding even though he was unable to keep the patron safe. Now suppose that, instead of Robert, Ocean Beach employs a lesser qualified lifeguard who finds himself in the same situation with a patron being pulled out by the rip current. If that lesser qualified lifeguard merely wades out into the water and tosses a lifebuoy in the general direction of the drowning patron without doing much of anything else, then he would fail meet his role responsibilities with respect to lifeguarding, and this would be so even if this particular lifeguard *cannot swim*, making it impossible for him to have executed a more worthy save attempt. Even if the lifeguard does not have the relevant kind of capacity needed for a good life-saving attempt, he still would have failed to "do his part" as a lifeguard.

Consequently, there is no automatic entailment from the fact that *S* has a certain kind of role responsibility to the fact that it is even practically possible for *S* to meet that role responsibility. There is no analogue of an "ought implies can" rule that applies to whether an agent has done one's part in the relevant way. Doing one's part as an occupier of a role might be a practically impossible task. Because of this, role responsibilities should be thought of as independent of a traditional kind of moral responsibility. It is possible for one to fail in meeting one's role responsibilities and to simultaneously deserve no moral blame (e.g., suppose that our lesser qualified lifeguard above, for whatever reason, was non-culpably thrown into the position of being a lifeguard) while it is also possible for one to succeed in meeting one's role responsibilities and to simultaneously deserve blame (e.g., consider a sentry in the

Wehrmacht who successfully alerts his superiors of an impending attack). What it means for an agent to meet one's role responsibilities is to perform well and to exercise a relevant capacity as an occupier of a role. Even if this is not a full-blooded deontological notion of responsibility, it can still be very useful in an evaluative practice. Someone who fails to do her part in this way is aptly criticizable as an occupier of the relevant role, even if doing so is not within her conscious control.

Returning now to the worry about doxastic involuntarism, even if it turns out to be true that we do not have the sort of conscious control over our beliefs that it would take to merit full-blooded deontological evaluation, one can still have inquiry-specific role responsibilities in the appropriate sense. One can still "do one's part" or not in inquiry and a belief can still exhibit a contributory capacity or not, even if we cannot come to our beliefs voluntarily in any relevant sort of way.[101] Suppose that an agent is simply incapable of forming a justified belief that p on the basis of the good evidence he has. Even if it is practically impossible for him to come to believe that p on the basis of his good evidence, he would still fail to meet his role-responsibilities and to "do his part" in gaining knowledge by doing anything other than believing that p on the basis of that evidence. If he did not to come to the appropriate belief, he would be aptly criticizable as an inquirer and would be so independently of considerations about doxastic voluntarism and voluntary control over belief.

---

[101] In this way, the account above has an advantage over the positions adopted by Monmarquet (1986), Steup (2000), Russell (2001), Ryan (2003) and others, who all see some kind of voluntarist thesis as crucial to the status of epistemically justified belief. Justification as a kind of role responsibility could, in a sense, be thought of as a deontological evaluation but the present view would not count as "deontology proper," as Nottelmann (2008) would put it. Though there are some important differences between the account outlined here and that of Feldman's (2001), the same could be said of his "role oughts."

Hence, what it means to meet one's role responsibility or to "do one's part" in this sense should not automatically be associated with exercising some kind of explicit, conscious control over a relevant action or state. "Doing one's part" is an agent's contributing toward an end in the appropriate way which, depending on the agent, may be an impossible task. For similar reasons, then, doing one's part in inquiry should also be distinguished from an epistemic requirement put forward by Roderick Chisholm: "We may assume that every person is subject to a purely intellectual requirement – that of trying his best to bring it about that, for every proposition h that he considers, he accepts h if and only if h is true." (1989, pg. 14) Excusing Chisholm's focus on true belief instead of knowledge, it still might be tempting to associate "doing one's part" in inquiry with this sort of analysis of a person's *efforts* in meeting the goal of inquiry.

"Doing one's part," however, is not merely to "do one's best" in the sense of putting forward effort. Our not-so-qualified lifeguard's life-saving attempt demonstrates this. Because the lifeguard cannot swim, he may very well have "done his best" in attempting to save the drowning patron by wading out into the water and tossing a lifebuoy, that being the best effort that he could have possibly put forward. What it is to "do one's best" in a given situation is constrained by the practical possibilities involved. Yet, even if that is the case, he has certainly not "done his part" as a lifeguard and has not contributed toward the goal of patron safety in a way that any decent lifeguard would have. His life-saving attempt does not exhibit the features necessary for this. In order to fully "do one's part" in achieving a goal that is constitutive of a certain kind of role, one must contribute toward that goal in a way that a good occupier of that role would, whether or not such a contribution is in the realm of practical possibility.

Serena's tennis-shot-making exhibition is also just simply not about how much effort she put into it or about whether or not she "did her best" in hitting the shot that she did. Excellent tennis players like Serena are often capable of making the shots they do without putting forward a lot of effort. Professionals like her are often capable of "doing their part" in hitting winning shots habitually in virtue of their finely-tuned automatic processing. We can even imagine an ideal kind of tennis player who, for her, tennis is so extraordinarily *easy* that tremendous efforts are simply not necessary in exhibiting the right kind of shot-making capacities. Even if tremendous efforts are not necessary for her, however, that would not thereby mean that she never "does her part" in the relevant sense. Considerations about effort and about the exhibition of the relevant capacities are separable.

Similarly, one can "do one's best" in inquiry and not have "done one's part" in the relevant sense. Suppose that Sam is inquiring about whether the earth is flat and he comes to be convinced that it is on the basis of, among other things, the fact that it *looks* flat and some conspiratorial considerations given to him by the Flat Earth Society. Now, it might be the case that Sam is just completely incapable of properly applying the ordinary kinds of deductive, inductive, and abductive considerations that would form a rational basis for the belief that the earth is round. Even so, he might very easily be "doing his best" in inquiring about this topic. Suppose that he is very conscientious about how he is forming his belief and that he has scoured all of the available sources of evidence that he can get his hands on. Even if this is the case, his belief still does not manifest the sort of capacity that any good inquirer would exhibit in inquiry. Perhaps, because of apophenia or some other abnormality in his psychology, Sam is disposed to

accept conspiracy theories in the absence of genuine support. Sam just does not have the proper constitution to deliberate about whether the earth is flat in a way that a good inquirer would. An inquirer who has "done his part" would not have believed on the basis of the considerations that Sam has, even though Sam "did his best."

The central claim on the table here is that there is a kind of derivative value relation, "doing one's part" in achieving a relevant goal, that makes being justified a good-making feature of beliefs within the critical domain framed by the constitutive goals involved in inquiry. Justified beliefs are epistemically better than unjustified beliefs even when these goals are not reached and "doing one's part" explains this. However, one might worry again that, since it is not the case that every belief is the product of an explicit inquiry, we have not given a complete explanation of the phenomenon in question. It is one thing to say, for example, that Anne's belief in Base Case exhibits her having "done her part" in inquiring about whether her husband is at home. The very nature of the case dictates that Anne is intentionally conducting an inquiry and that she actually has the psychologically instantiated goals that are constitutive thereof. It is another thing to say that *every* justified belief exhibits an agent's "doing her part" in inquiry, especially those that are formed and maintained unintentionally and below the level of consciousness. Yet, the epistemic goodness of such beliefs purportedly relies on such a status. Can we make sense of this?

Recall that, in critical domain evaluation, we evaluate *from a certain perspective*. In the case of epistemic evaluation, we evaluate from the perspective of an inquirer or someone interested in the constitutive goals of inquiry, independently of whether the item being evaluated (e.g., a belief) belongs to an actual inquirer (see pgs. 38-40 in Ch. 1). If

someone unintentionally forms and maintains a belief that *p* below the level of conscious awareness and is not at all inquiring about whether *p*, that belief will still be evaluable from the perspective of an inquirer. That belief still might exhibit the good-making features associated with the office of inquiry or it might not. The agent who has the belief may have formed it in a way that she has contributed toward knowledge, basing it on the right sort of considerations, and doing her part from the perspective of an inquirer. Consciously formed with the explicit goal of inquiry in mind or not, her belief still might exhibit the relevant sort of capacity that it would take for it to be good from the perspective of that evaluative practice.[102]

In this way, the evaluative practice framed by the constitutive goals of inquiry is somewhat unique in how widely applicable it is. Many of our beliefs are formed and maintained unreflectively and independently of our more explicit inquiries. Since inquirers often purposefully end up forming beliefs about various topics, beliefs are a target of evaluation within that critical domain. The evaluative practice then extends out beyond our explicit inquiries, applying in many circumstances to people who do not actually occupy the role of being an inquirer (but yet form and maintain beliefs). Contrast this with, say, the critical domain framed by the goals involved in competitive tennis playing. Competitive tennis players purposefully hit tennis shots, making such things evaluable from that perspective. It is not the case, however, that the vast majority people in the world, most of whom are not actually tennis players and are not even

---

[102] This is one difference between the kind of framework suggested here and the sort of framework that Foley (1987) lays out for "epistemic rationality." In the way that Foley sets things up, the epistemic rationality of a belief (at least in its most primary sense) depends on an agent *actually having* the goal of now having true beliefs and now not having false beliefs and whether the belief would be reflectively thought of as a non-trivial means to that goal. See, e.g., the way he describes his position on pgs. 8-12. He does eventually mention a more "secondary" form of epistemic evaluation whereby we can evaluate the beliefs of an agent "as if" she had the goal, but this is clearly not his focus.

interested in tennis playing, go around unreflectively hitting tennis shots, opening themselves up to evaluation from that perspective. Unlike the formation and maintenance of beliefs, it does not really even seem possible for someone to play tennis and hit tennis shots in a way that they are *completely* unaware that they are doing so. Because beliefs do have that kind of status, people are often aptly evaluable and criticizable from the epistemic perspective whether they like it or not.

Consider another worry that one might have in reflecting on Base Case. The idea in that case is that Anne's belief is a good one because it exhibits her having done her part in achieving the goal of knowledge; it exhibits her having met her role responsibilities with respect to the office of inquiry even when her belief turns out to be incorrect. But, has she really "done her part" in that way? Surely enough, the basis on which she forms her belief does seem to adequately support believing that her husband is home. Yet, as an inquirer, she could do more to ensure that she gains knowledge about the topic that she is interested in. For instance, she could simply walk into the room where she thinks her husband is sitting and attempt to visually verify that it is, indeed, him. Abstracting away from evaluative interests other than those involved in successful inquiry, it seems as though, ideally, Anne would contribute toward the goal in that way. Perhaps this should make us question whether her belief really has exhibited a meeting of the relevant role responsibilities. Insofar as we still think that Anne's belief is epistemically justified even if she has not attempted this visual verification, we should also question whether "doing one's part" and "meeting one's role responsibilities" can explain the epistemic goodness of such a status.

To handle this worry, a distinction needs to be made within the epistemic evaluative practice between what I will call one's "doxastic" and "evidence-gathering" contributions. There are different ways that an agent may make a contribution toward the ideal end of inquiry. For instance, when an individual believes justifiably that p, she makes a particularly *doxastic* contribution toward the goal with respect to p. She takes the appropriate doxastic attitude concerning p by believing it on the basis that she does, contributing in the moment toward knowledge. We can think of satisfying the related role responsibility as meeting a kind of conditional demand indexed to a particular time (call it "dox-RR"): If S has a level of support such that she is propositionally justified in believing that p at time t, then S must believe that p on that basis at t. In other words, an agent must capitalize on her propositional justification. Of course, whether an agent must believe that p at a particular time will depend on the kind of justificatory support available to her. Hence, believing that p on some basis may be an inquiry-related role responsibility at one time but not at another.

However, even if someone meets dox-RR by justifiably believing that p at a particular time, this does not necessarily mean that there are no other ways in which she could contribute toward the ideal satisfaction of the goal of inquiry, even with respect to p alone. Generally speaking, unless one's justificatory support in a given moment completely clinches whether p, an agent can contribute further toward the ideal satisfaction of the goal by *gathering additional evidence* with respect to whether p. For instance, even though Anne has adequate justificatory support for believing what she does about whether her husband is at home (making her belief epistemically good), her epistemic situation might be improved if she were to gain additional evidence. In her

case, she might learn that her belief is false. Until the matter is completely clinched with respect to whether p or until there are no more untapped evidential resources with respect to whether p, there are always more ways in which an agent can contribute in inquiry. Of course, an inquirer's contribution by gathering more evidence is not a particularly *doxastic* kind of contribution, but it is still a consideration worth highlighting within the critical domain.

In Base Case, let time t_1 be the moment in which Anne gathers the last bit of evidence about whether her husband is at home (i.e., when she hears the munching noises in front of the television) and thereby forms the belief that he is. The point here is that, at t_1, Anne discharges a particular kind of role responsibility (i.e., dox-RR) by justifiably believing that her husband is at home. As suggested above, this makes it the case that her belief has epistemic value. Yet, since it is not completely clinched as to whether her husband is home, she could still contribute more to the ideal satisfaction of the goal of inquiry. She could still gather more evidence after t_1. When she does not do so, she is, at the very least, less than ideal in her evidence gathering and her overall, domain-internal contributions. Even so, it is still an epistemically good thing for her to have contributed in the way that she did and for her belief to exhibit the satisfaction of dox-RR.

It is important to be able to explain *both* the epistemic goodness of justified belief and the epistemic goodness of evidence gathering. Any proposed epistemic axiology must be able to explain the basis for these distinct values if it is to be an adequate one. This is not necessarily a simple task. Take, for instance, a problem that has been raised about formal epistemological frameworks based on Bayesian probability theory. Among other things, Bayesians have been accused of not being able to account for why we

should want to conduct new experiments and gather new evidence about the world if our beliefs already happen to fall in line with the probability calculus.[103] If the object of our distinctively epistemic desires is merely to achieve probabilistic coherence and to conditionalize when new evidence comes in, then, insofar as our current beliefs already meet those constraints (making them Bayesian-justified), we may be best off resting content with our beliefs as they are, refraining from experimentation and ignoring as much new information as possible, instead of risking a disruption to our belief set. Now, whether or not Bayesians have satisfactorily handled such worries, there is an important recognition here that it is not good enough to merely tell us that justified beliefs are epistemically valuable. An adequate epistemic axiology must also explain why it is good for us to go out into the world and gather new information, to experiment, and to beef up our evidential base (even when we already have justified beliefs). An axiology based on the goals involved in inquiry can do this in a unified way.

So, it is epistemically good for an agent to meet the demand in dox-RR even though doing so is not the only way that an agent might contribute in inquiry. However, it seems as though dox-RR must be just part of a wider kind of doxastic demand within the epistemic evaluative practice. Consider the following scenario: Suppose that, before time t_1, A and B are "propositional justification twins," possessing the exact same amount of justificatory support for each and every proposition. Further, suppose that, before t_1, A and B hold the exact same beliefs about the world on the exact same bases. At t_1, however, while neither A nor B has any evidence with respect to whether p, A comes to believe that p unjustifiably while B withholds belief about whether p. As a consequence,

---

[103] For an overview of the worry here, see Good (1967), Skyrms (1990), and Oddie (1997).

at t_1, A and B have the exact same amount of justified beliefs about the world but A has

one more unjustified belief than B. Our intuitive response to a scenario like this is that,

from the epistemic perspective, A's belief set is worse off than B's. More, it seems to us

that A's is worse off than B's because, with respect to p, A fails to meet a kind of doxastic

demand from the epistemic perspective that B respects. Yet, if meeting dox-RR were the

only salient kind of doxastic demand, we would be left unable to explain this; A and B are

equal with respect to meeting dox-RR.

The fact of the matter is that, in relation to some particular proposition p, the

doxastic demand on an agent from the epistemic perspective requires not only believing

that p when enough justificatory support is available, but also withholding or suspending

belief when it is not.[104] One must proportion one's doxastic stances to the support

possessed, so to speak. If one fails to withhold belief when there is not enough

justificatory support available (thereby forming an unjustified belief), then the belief that

has been formed is an epistemically bad one. While this claim would explain the

difference in epistemic value between the belief sets of A and B above, we need to

provide some sort of basis for it. All along we have been discussing the positive

fundamental and derivative values arising from the critical domain framed by the

constitutive goals of inquiry. We have not yet discussed any fundamental or derivative

---

[104] Of course, believing that p or withholding belief about whether p does not exhaust the logical space of doxastic stances one might take toward p. For instance, one might also *disbelieve* that p. Here, I will be assuming that disbelieving that p is the very same thing as believing that not-p. In this way, I will not be focusing on it as a doxastic stance in need of theorizing over and above what has already been said. Justifiably disbelieving that p is epistemically good in the same way that justifiably believing that p is; it exhibits an agent's doing one's part in satisfying the constitutive goals of inquiry about whether p. I will also not be focusing on any doxastic stances related to what have been called "degrees of belief" or "credences," though nothing I say below will rule out quantitative frameworks if one prefers that way of thinking about belief and justificatory support.

disvalues within such a domain. This is what it would take in order to properly explain the comparative verdict.

When one forms a justified belief, one meets a certain role responsibility vis-à-vis inquiry, contributing toward knowledge in a certain way. However, when one forms an unjustified belief, one seemingly *fails* to meet an inquiry-related role responsibility. At the very least, it seems that a good inquirer would withhold belief when there is not enough support available. This is a genuinely *bad* thing from the epistemic perspective. As an analogy, take the lifeguard who merely tosses a lifebuoy out to a drowning patron. Whether or not this life-saving attempt is successful, the lifeguard fails to contribute to the situation in a way that a good lifeguard would, failing to meet a role responsibility. Because the lifeguard fails in this way, his life-saving attempt is not merely evaluatively neutral from the perspective of the lifeguarding evaluative practice; it is a genuinely *bad* life-saving attempt.

What, though, could explain this evaluative verdict? The lifeguard's failure is bad because there are certain fundamental *disvalues* that (at least partially) govern the lifeguarding critical domain. From the perspective of lifeguarding, certain kinds of physical harms suffered by patrons are genuinely disvaluable things. The lifeguard fails with respect to a role responsibility because he fails to contribute properly to the *avoidance* of that kind of harm. As such, the life-saving attempt is *derivatively* disvaluable as it manifests such a failure. Regardless of whether or not any patrons are actually harmed, it is bad because the lifeguard fails to do his part in its avoidance. Yet, the badness of the attempt here is ontologically dependent upon the more fundamental badness of patron harm. If it were not the case that certain kinds of patron harm were

disvaluable, then the lifeguard's failing to do his part in the avoidance of that kind of patron harm would also not be disvaluable.

If we want this comparison to be instructive about the epistemic evaluative practice, then we should attempt to incorporate a similar kind of negative valuation. What, though, might play the role that patron harm is playing with respect to lifeguarding? What are the most fundamental disvalues within the epistemic critical domain? Recall the suggestion that the most fundamental *positive* values within the domain are the ideal ends of inquiry and, in the case of propositional inquiry, knowledge. If instances of knowing whether *p*, for any *p*, are fundamentally valuable, it might be natural to think that instances of *ignorance* about whether *p* (or the lack of knowledge about whether *p*) have a fundamental *disvalue*. Yet, this cannot be correct. From the epistemic perspective, it is not necessarily a disvaluable thing to be ignorant given certain circumstances. For instance, if I am ignorant about whether there are an even or odd number of stars in the universe, that fact, by itself, does not generate epistemic disvalue. My not knowing whether there are an even or odd number stars is instead some kind of epistemic evaluative limbo. It is certainly not an epistemically *good* thing for me or anyone else to be ignorant, and it is certainly less than epistemically ideal to not know things, but ignorance is not an epistemically disvaluable thing in and of itself.

Consider an example to make this point a bit clearer: Suppose that I am a competitive tennis player and that my match today has been postponed until tomorrow because of the current rainy conditions. From the evaluative perspective of competitive tennis playing, while it is certainly not a *good* thing that my match has been postponed, and while it is certainly not the *ideal* state of affairs, the postponement is still something

that sits in evaluative limbo. It is not like I played the match and lost, which *would* be a

kind of fundamental competitive tennis playing disvalue. Rather, the postponement is

just an evaluatively neutral state of affairs and, *ceteris paribus*, we do not really see any

real value or disvalue until the match gets resolved. Something similar holds for the state

of ignorance with respect to epistemic evaluation. Simply not knowing whether *p* is

much more like the postponement of a tennis match than it is like playing a tennis match

and *losing* it.

For a more promising suggestion about fundamental epistemic disvalue, we might

take a hint from William James:

> There are two ways of looking at our duty in the matter of opinions, - ways
> entirely different, and yet ways about whose difference the theory of knowledge
> seems hitherto to have shown very little concern. We must know the truth; and
> we must avoid error, - these are our first and great commandments as would-be-
> knowers; but they are not two ways of stating an identical commandment, they are
> separable laws. (1956, pg. 17)

Now, James's focus here is explicitly on "duty in the matter of opinions" and on

"commandments" and "laws" related to our beliefs, not necessarily on the fundamental

values and *disvalues* within the epistemic domain. Nonetheless, I think that his

suggestion can be helpful in the current context. The idea seems to be that, in

epistemology, we have both things to bring about and things to avoid. More relevantly

for the current discussion, we have both positive and negative goals and, by association,

fundamental values and disvalues.

According to James, it seems that the most fundamental values in the epistemic

domain would be instances of knowledge. By contrast, the most fundamental disvalues

would be instances of *error* or *false belief*.[105] If this is correct, it would suggest a story

about unjustified belief that is analogous to the kind of bad life-saving attempt already

discussed: Unjustified beliefs are derivatively disvaluable as they manifest a failure by

the agent to do her part in avoiding error, which is a fundamental disvalue. Take the case

above in which *A* forms an unjustified belief on the basis of no evidence at all. By failing

to withhold belief, *A* fails to "do her part" in avoiding error regardless of whether her

belief turns out to be true or false. She fails to exhibit a relevant kind of capacity in

contributing toward the avoidance of error, a kind of doxastic sensitivity that a good

inquirer would exhibit. As such, she is aptly criticizable from such a perspective.

Generally speaking, unjustified beliefs manifest this kind of failure, and that is the case

even if the relevant believing agents luckily do not actually end up in error. Failing in

such a way makes the associated beliefs epistemically disvaluable. Yet, the manifestation

of such a failure is derivatively disvaluable relative to the more fundamental disvalue of

error. If error were not epistemically disvaluable, then neither would be the failure to

contribute properly to the avoidance of it.

  The entire framework presented in these chapters relies on there being a

connection between epistemic evaluation and the goals involved in inquiry. Should we

really think that error is a genuine *disvalue* from the inquirer's perspective? This may

---

[105] This is certainly not the first time that James's quote above has been appealed to in discussing epistemic norms and values. Interestingly, however, most who do bring up James's dual commandments ignore the fact that he claims that it is *knowledge* that is the positive goal in epistemology and then slide to claims about true belief being the goal (see, e.g., Riggs 2003 and Piller 2009). There is some reason for this. James is systemically loose in his discussion and he seems to be perfectly content in switching back and forth between claims about knowledge and claims about true belief (see, e.g., the way in which he further describes his points on pg. 18). The suggestion in these chapters is that the knowledge claims are more appropriate. Tracing the roots of James's dual commandments even further back, the Cartesian project in the *Meditations* can also easily be seen as endorsing a kind of view in which knowledge (in the sense of incorrigibility) has a fundamentally valuable status while error has disvalue.

seem especially odd given that we have previously denied that true belief is fundamentally valuable. It should seem plausible that error is genuinely bad from the epistemic perspective, independently of the value of true belief, because it actively *thwarts* the satisfaction of the constitutive goals of inquiry. For instance, insofar as one desires to know whether *p*, being in error about whether *p* would prevent the state of affairs in which that desire is satisfied from ever obtaining. Moreover, error about whether *p* would thwart the goals of *other* inquiries one might engage in insofar as other beliefs are formed on it as a basis. It would make sense from the inquirer's perspective that error is seen as a kind of fundamental enemy. Hence, there is still a sense in which the old slogan is true: epistemically-minded agents try to gain truths (i.e., to know truths) and avoid falsehoods.

To close this chapter, I will discuss a potential worry from Selim Berker's (2013a) extended discussion on what he calls "epistemic teleology."[106] The approach taken in these chapters can appropriately be thought of as an instance of the kind of teleological theory that Berker argues against. The theory outlined above explains what we are doing in our epistemic evaluative practice in terms of inquiry which is a goal-directed activity. Because inquiry constitutively aims at knowledge, I have claimed that knowledge is a fundamental epistemic value or, as Berker would call it, a relevant "final value." In turn, other items take on epistemic value in virtue of their derivative value relations, both instrumental and non-instrumental, to knowledge, constituting what Berker would call "a theory of overall value." These claims have culminated in this chapter in a story about the place of epistemic justification. If I am correct, our evaluations about epistemic

---

[106] See, also, Berker (2013b) and the exchange between Berker (2015) and Goldman (2015).

justification are explained by appealing to these values (and disvalues) associated with the goal-directed practice of inquiry.[107]

Berker's worry is that such a teleological theory, one that specifies the role of epistemic justification in terms of its place within a value structure, will necessarily generate unpalatable tradeoffs in identifying what we would be justified in believing. According to Berker, there is a general problem that value-first teleological theories face given their propensity to sanction sacrifices that we should not be willing to accept. As an analogy, Berker (pg. 358) has us consider some of the difficulties faced by act utilitarians. Because they identify what is *morally justified* in terms of how an action relates to a certain value structure (whether the theory be maximizing, satisficing, or require some other relation to that value set), they have difficulties in dealing with cases where something valuable might be sacrificed in order to achieve what would be seen as a greater good, but doing so in an intuitively unsavory way.[108] Take, for instance, Thomson's (1976, pg. 206) case about a doctor who has the opportunity to sacrifice the life of a healthy patient by killing her and harvesting her organs in order to save the lives of five terminally ill patients. The difficulty for the utilitarian is that, because she identifies justified action in terms of how it relates to an overall value structure, she will have a hard time condemning what would otherwise be seen as an unjustified action because the action in question generates a great amount of value (perhaps even the most value of the actions available).

---

[107] Because the theory presented here focuses both on teleology and virtue, it is best construed as a version of what David Solomon (2003) calls "routine" virtue epistemology (as opposed to "radical"), one that affords a relatively basic position to virtue in the theory while simultaneously keeping intact a commitment to teleological explanation.

[108] Or, as Rawls (1963, pg. 124) would put it, in a way that ignores the "separateness of persons."

By bringing up such cases, Berker's point is not to argue against act utilitarianism or any other particular teleological ethical theory. From his perspective, such ethicists may or may not be able to deal adequately with such tradeoff cases. What he does attempt to do, however, is illustrate the inherent propensity of teleological theories for generating these sorts of questions about the tradeoff of value. Independently of whether certain ethical theories can properly deal with the kinds of tradeoff questions generated by their endorsed value structures, Berker argues that teleological *epistemic* theories cannot. According to him, teleological epistemic theories will always generate problematic tradeoffs and will always flout what he calls the "separateness of propositions."

He begins to illustrate this point by drawing on an example from Richard Fumerton (2001, pg. 55) which I will slightly modify here in order to make the issue more salient: Suppose that I am a scientist who is seeking a large research grant from a religious institution who only awards such grants to theists. Also suppose that, after careful examination of the available evidence, I accurately discern that the evidence does not support believing that God exists. However, I know that if I were somehow able to convince myself that God exists, I would receive the grant and the subsequent research money would be a causal means to gaining a lot of knowledge about the world that I otherwise would not gain. The question is now the following: Do these facts make it the case that I would be *epistemically justified* in believing that God exists? Clearly, the answer must be "no" and, yet, as Berker claims, there would be a certain kind of teleological epistemic theory that would not be able to generate this verdict. Consider a teleological theory that takes knowledge to have a fundamental epistemic value and then

identifies justified belief in terms of what maximizes overall epistemic value given the situation. It would then seem to be the case that I would be epistemically justified in believing that God exists. Even though I would be sacrificing a certain kind of epistemic value with respect to one proposition, I gain much more with respect to others. Because such a theory does not respect the "separateness of propositions," it is untenable.

As Berker acknowledges, the Fumerton case will only plague the most primitive kind of teleological theory and it cannot be counted as an objection to any existing theory in particular. It does, however, signal to him a more general kind of problem. Berker claims that, no matter what the teleologist's fundamental values are, no matter what kinds of derivative value relations the teleologist allows in her theory of overall epistemic value, and no matter what kinds of principles are appealed to in linking such values to epistemic justification, the resulting theory will be indefensible. Berker spends many pages making cases against a variety of adjustments the teleologist might want to make in order to avoid problematic inter-propositional tradeoffs, fielding theories that restrict derivative value relations so as to exclude nonproximate causal means, theories that exclude causal means altogether, theories that time-relativize fundamental value, theories that agent-relativize it, proposition-relativize it, etc.[109] In general, Berker argues that, no matter what adjustments are made, teleologists will always allow for the construction of problematic tradeoff cases in a way that the "separateness of propositions" is not respected or they will make *ad hoc* restrictions to their theory in a way that makes them not worth pursuing (or both).

---

[109] Particularly interesting is Berker's discussion of Foley's (1993) time-relativized teleological solution.

Given this, it will be worth examining whether the teleological theory presented in these chapters is guilty of the charges that Berker imagines it to be. Recall that the claims in this chapter are about the domain-internal value of epistemic justification from the perspective of an evaluative practice that takes the constitutive goals of inquiry to be fundamentally valuable. The most central claim is that an epistemically justified belief is derivatively valuable as it manifests a kind of agential contribution toward knowledge of a particular proposition. When a belief about whether p is formed in a way that it enjoys the requisite amount of support, an agent "does her part" for the moment in gaining knowledge about whether p, exhibiting an inquiry-relevant virtue. Hence, the value of epistemically justified belief is explained teleologically through the goal-directed practice of propositional inquiry.

If this story is the right one, then evaluations about epistemic justification are indexed to a particular propositional inquiry type. It is true that, within the critical domain outlined here, knowledge of any proposition whatsoever has a fundamental epistemic value (see pg. 34 in Ch. 1). This will, as Berker imagines, generate certain questions about tradeoffs in value and it allows us to weigh various states of affairs against one another for their overall epistemic worth. Yet, in our evaluations about whether or not a particular belief is epistemically justified, we contextually restrict the domain of fundamental epistemic value to the constitutive goal of a particular propositional inquiry, ignoring the implications that the uptake of a relevant belief would have on the outcomes of other, unrelated inquiries that one might engage in. When we ask about whether a belief that p is epistemically justified, it is the constitutive goal involved in inquiring about whether p that is relevant and nothing else. A belief about

whether *p* is epistemically justified, then, when the agent can count as having "done her part" in an inquiry *about whether p*, not necessarily when she has maximized overall epistemic value or has met a certain satisficing level of epistemic value or has met any other principle tying the uptake of her belief to epistemic values unrelated to the constitutive goal of that particular propositional inquiry.

This kind of contextual restriction of value within critical domains is not unique to epistemic evaluation. For example, we willingly make such restrictions in our evaluations of "good shots" in competitive tennis playing as well. Consider, once again, The Shot (see pgs. 113-4). Suppose further that, because Serena realizes that she has hit a good shot, she becomes overconfident in her abilities and that this causes her to lose many more tennis matches in the future than she otherwise would. The fact that her shot generates a significant amount of domain-internal disvalue does not make us less inclined to say that the shot was still a good one, that it is still constitutively good relative to the narrower goal of the shot itself. In the moment, she still contributes in a way that a good tennis player would toward hitting a winning shot, manifesting a competitive tennis playing virtue. In our evaluation of whether or not her shot was a good one, manifesting such a virtue has a very important standing to us. When we say that it was a good shot, we are willing to "put blinders on," so to speak, as to how that particular shot affects Serena's other shots in other, independent points and matches. This does not make it any less the case that we can explain why such a shot has the good standing it does *teleologically* relative to the goals involved in competitive tennis playing. It is still good derivatively, relative to the goodness of winning tennis matches and hitting winning

tennis shots. The claim here is that we do something very similar in our evaluations about epistemic justification.[110]

So, what about Berker's objections? Will the proposal *still* allow for the construction of problematic tradeoff cases? I do not think so. Consider, again, the Fumerton case. For simplicity, let us suppose that the uptake of the belief that God exists in that case would generate 1,000 extra bits of propositional knowledge that I otherwise would not gain. Two relevant evaluative verdicts can then be rendered from the epistemic perspective: (1) that my belief that God exists is epistemically unjustified because I have not properly "done my part" vis-à-vis inquiry about whether God exists and (2) that the overall state of affairs in which I unjustifiably believe that God exists (and also gain 1,000 extra bits of propositional knowledge) is *epistemically better* than the relevant alternative overall state of affairs in which I justifiably do not believe that God exists. Do verdicts like these indicate some sort of tradeoff problem for the view? I do not see how they could. They seem to be completely appropriate. There is something like a tradeoff of epistemic value being endorsed in the second verdict but it is endorsed in a way that is independent of the question about epistemic justification. Given the proposal above, there is no reason to think that verdicts about epistemic justification entail anything about the overall epistemic goodness of the resultant state of affairs. Epistemic justification will not necessarily maximize epistemic value nor will it

---

[110] To be clear, the kind of proposition-specific, contextual restriction proposed here is significantly different than the "proposition-relative" solution explicitly discussed by Berker. The proposition-relative solution that Berker does discuss suggests individuating *kinds* of epistemic value relative to each proposition, where our epistemic goals have, as a part of their content, an index to a specific proposition. If we did individuate our epistemic goals in this way, then particular states of affairs could not have epistemic value *simpliciter*, but could only have *p*-epistemic value or *q*-epistemic value or *r*-epistemic value, etc. However, to contextually restrict our epistemic evaluations in the way suggested here, we need not appeal to such individuation. Berker does not explicitly consider a solution in this vein.

necessarily meet some satisficing threshold (e.g., consider a case in which a justified belief leads to many false beliefs).[111] Yet, the theory in question does seem to be teleological, one in which epistemic justification is a kind of derivative, domain-internal value.

Is the proposal here objectionably *ad hoc*? What is doing the work in avoiding tradeoff problems is the claim that, in our evaluations about epistemic justification, we contextually restrict the relevant domain of fundamental value to the ideal end of a single propositional inquiry. If we *stipulate* that this is the way our evaluations about epistemic justification work, then, of course, that may automatically rule out problems that other kinds of teleologists might have in respecting the "separateness of propositions." Is such a stipulation something we should be worried about? I do not think that it is. The kind of stipulation in question is *ad hoc* in the same way that any evaluative practice might be somewhat *ad hoc*. Contextually restricting the domain of value just seems to be something we *do* in our epistemic evaluative practice and that is just the result of it being *our* practice. The epistemological critical domain might easily exhibit some sensitivity to our own, collective interests and, somewhere along the line, we have taken a special interest in the kind of inquiry-relevant virtue described above. It is that particular focus that explains why our evaluations about epistemic justification are the way they are, being a kind of evaluation that ignores other contextually irrelevant propositions.

---

[111] Michael Slote has pointed out that this approach to epistemic justification has somewhat of a Kantian flavor. For Kant, agents who act in morally justified ways treat each individual as an end and not as a mere means to other sources of value (moral or otherwise). Similarly, since we index the relevant evaluations to particular propositions, agents with epistemically justified beliefs treat the constitutive goal of a particular inquiry as an end in itself and not as a mere means to other sources of epistemic value.

This kind of reply might point someone like Berker in the direction of a bigger issue with the teleological theory presented above. If the theory is correct, Berker might rightfully wonder what the overall significance of epistemic justification is supposed to be. For instance, given the place that evaluations about epistemic justification have within the epistemological critical domain, it does not seem as though verdicts about justification are going to definitively tell us anything about what we *should* believe (or what we *ought to* believe) in a given situation. Justification's place within the domain seems too gerrymandered to give us any clear answers, by itself, about overall prescriptive force on belief. It merely tells us whether an agent's belief exhibits a circumscribed kind of domain-internal virtue.

For people like Berker, though, telling us what we *should* believe is exactly what verdicts about epistemic justification are supposed to do:

> As I see it, the most fundamental question in ethics is 'What should I do?' The exact wording of this question can take different forms, depending on one's normative vocabulary of choice ('What do I have most reason to do?' 'What is it rational for me to do?' 'What am I justified in doing?')… I think it is useful to view epistemology as having a parallel structure. On this way of conceiving of the discipline, the most fundamental question in epistemology is 'What should I believe?' The exact wording of this question can take different forms, depending on one's normative vocabulary of choice ('What do I have most reason to believe?' 'What is it rational for me to believe?' 'What am I justified in believing?'). (pgs. 337-8)

Berker's conception of the evaluative project in epistemology is analogous to a kind of Moorean conception of ethical evaluation where the most basic questions fall within a kind of unqualified normative or prescriptive domain.[112] According to these conceptions, correct verdicts about morally justified action and about epistemically justified belief will always have genuine prescriptive force in virtue of the kinds of evaluations that they are.

---

[112] See Ch. 1 in Moore (1903).

They will automatically tell us what we *ought* to do and what we *ought* to believe, respectively. However, it is a substantive assumption that this is what epistemologists are up to and, among other things, it seems to rule out the possibility that epistemic evaluation is an instance of a critical domain in the way described above. For many reasons, the relationships among the correct epistemic evaluative verdicts and genuinely normative prescriptions on belief are a complex matter. This will be the topic of the final chapter.

Chapter 4: What We Ought to Believe

This chapter is devoted to defending the view outlined above, that epistemic evaluation is a kind of critical domain framed by the constitutive goals of inquiry, against what might be otherwise taken as an objection from considerations about the ethics of belief. Evidentialists (or "epistemic exclusivists") in the ethics of belief think that epistemic justification and epistemic rationality play a unique and exclusive role in determining the normative facts about our beliefs. At least *prima facie*, one might think that the truth of evidentialism would conflict with the framing of epistemic evaluation as a kind of critical domain as the nature of critical domains leaves such a status unexplained. In the first section of the chapter, I clarify this issue and make the relevant explanatory desiderata more transparent. In the second section, I argue that certain defenses of epistemic exclusivism (i.e., those utilizing what have been called "no rewards principles") are bound to fail in challenging the critical domain picture above because such arguments actually imply the falsity of certain forms of evidentialism. Finally, I argue that the remaining lines to evidentialism are either dialectically ineffective or are consistent with the critical domain picture suggested above. I discuss an argument involving the so-called "wrong kinds of reasons," eventually making a distinction that can insulate the view here against attacks issuing from the ethics of belief.

4.1: The Simple Connection Thesis, Evidentialism, and Guidance-Giving Oughts

In the previous chapter, Selim Berker's (2013a) objections to teleological views of epistemic justification were discussed. In due course, it was brought out that Berker

views correct verdicts about epistemic justification as definitively settling the question of what we ought to believe in a similar way that verdicts about moral justification (purportedly) settle the question of what we ought to do. Of course, Berker is aware that analogies between the epistemic and moral realms are bound to break down at some point but, as he would say, "that's what makes them analogies and not identities." (pg. 3) Nevertheless, one of his central principles for theorizing about epistemic justification is that, like evaluations about moral justification, there is a special kind of authoritative and prescriptive normative force that follows from correct evaluations about it.

Because of this, Berker commits himself to something like the following view: One ought to believe in accordance with one's epistemic justification. There are a number of ways to further precisify a claim like this. We might, for example, take it as a conjunction of three distinct "ought" claims (call this conjunction "The Simple Connection Thesis" or "SCT"): For any agent S and any proposition p, (1) if S is propositionally justified in believing that p, then S ought to believe that p, (2) if S is propositionally justified in believing that $\neg p$, then S ought to disbelieve that p (or believe that $\neg p$), and (3) if S is neither propositionally justified in believing that p nor propositionally justified in believing that $\neg p$, then S ought to suspend belief as to whether p. As such, SCT utilizes a narrow-scope "ought" operator in three distinct conditionals which, together, are often seen to exhaust the logical space of doxastic attitudes one might take toward a given proposition.[113]

---

[113] Whether there is an exception to the exhaustion claim here depends on the nature of doxastic states like belief and whether or not, e.g., what falls under David Christensen's (2004) "graded conception" of belief is a genuine phenomenon. If credences are the best way to think about our doxastic attitudes, then even though SCT would need some precisifying, there should still be a relevant way of formulating the kind of normative connection Berker has in mind between a belief's epistemic support and what we ought to believe. Consider a first pass: We ought to believe propositions to degrees that are proportional to the amount of epistemic support they have.

Berker never explicitly endorses SCT or any other precise version of the connection between epistemic justification and what we ought to believe. Yet, given what he does explicitly say, it would be fair to assume that he thinks that *some* kind of special and necessary connection like the one in SCT is appropriate. After all, in the passage quoted at the end of the previous chapter, it seems as though Berker thinks that asking what we ought to believe is, in some sense, the *very same question* as asking what we are epistemically justified in believing (see pg. 149 in Chapter 3). He is not alone in thinking that there is a very tight connection. A cluster of related views can be found in the debate between the so-called "evidentialists" and "pragmatists" in the ethics of belief literature. Like Berker, evidentialists think that epistemic considerations are ultimately all that matter in determining what we ought to believe and that all other kinds of considerations (e.g., the purely practical benefits of believing) can be ignored. In other words, evidentialists think that epistemic considerations are *exclusive* in establishing the normative facts about belief.[114]

There are a number of ways to carve out different evidentialist positions, but anyone who falls under the evidentialist label in the ethics of belief takes up some kind of exclusivist stance.[115] It is a bit of shame that these exclusivists are called "evidentialists." Not all positive epistemic considerations or reasons are clearly evidential in character.

---

[114] See Steglich-Petersen (2009) for a seminal source of the "exclusivity" terminology.

[115] One major dispute among contemporary evidentialists is related to the question of *why* epistemic considerations are exclusive in determining what we ought to believe. For example, traditional evidentialists like William Clifford (1947) seem to think that epistemic considerations are exclusive for *moral* reasons, because of the moral harms that would be incurred if we were to believe otherwise. As will be shown, however, an evidentialist need not use the requirements of morality at all in attempting to provide a basis for epistemic exclusivity. It should also be noted that the above way of framing Clifford's views relies on a certain charitable reading of his essay as neither he nor William James (1956) in their disagreement make a clear distinction between moral evaluation and any other kind of evaluation. See Haack (2001) for more on this.

Consider, for example, the kind of epistemic support relations that hold between considerations about parsimony and a particular set of theories. Are such considerations properly thought of as *evidence* for one theory over another? What about the considerations that justify believing in the principle of non-contradiction or in the validity of mathematical induction? Are those considerations *evidence*? Perhaps, but only if "evidence" just means something like "that which contributes to epistemic justification."[116]

On a more natural reading of "evidence," it would be entirely consistent for someone to hold that epistemic considerations are exclusive in determining what we ought to believe while simultaneously denying that all epistemic considerations count as evidence. Similarly, those who hold a kind of negative theory about epistemic justification (see footnote 94 in Chapter 3) sometimes do not require any positive epistemic reasons at all in order for a belief to be justified. Again, it would be entirely consistent for someone like that to accept, say, SCT and to simultaneously deny that *evidence* is doing all of the normative work. In this chapter, however, I will follow tradition and will use the term "evidentialism" to refer to exclusivism about epistemic considerations. The above reservations should still at least be noted.[117] Pragmatists, then, are those who, given that there are normative facts about our doxastic attitudes, deny evidentialism and give other, non-epistemic considerations "a seat at the table," as it were.[118]

---

[116] See Kelly (2014) for more on the relevant ambiguities with "evidence."

[117] I am certainly not the first to note the narrowness of the "evidentialist" label. Also see, e.g., Adler (2002, pg. 5), Schroeder (2012), and Berker (unpublished manuscript). For an array of evidentialist positions, see Clifford (1947), Railton (1994), Kelly (2002), Shah (2006), Thomson (2008), Parfit (2011), Whiting (2014), and Nolfi (forthcoming).

[118] For some pragmatist positions, see James (1956), Reisner (2009; forthcoming), Markovits (2014), Rinard (2015; 2017), and Leary (2016).

Evidentialism is relevant for the following reason: If evidentialism is true, then the status of epistemic justification (or perhaps epistemic rationality) plays a unique and crucial role in determining the normative facts about belief. Most importantly, epistemic justification would be unique in a way that the circumscribed value relation described in the previous chapter might not be able to account for. Why, for example, should we think that a contextually restricted, domain-internal evaluation about whether an agent has "done her part" in inquiry plays an *exclusive* role in determining whether we ought to believe something? At least *prima facie*, the pragmatist position in the ethics of belief is better suited for the view. Whether or not one would be doing one's part in inquiry by believing a certain way only seems to be one kind of consideration among others (albeit an important one) in determining the normative facts about what doxastic states we ought to take up (if there are such facts at all). If epistemic justification is domain-internally valuable in the way I have claimed, why should all other kinds of considerations (e.g., other, non-inquiry-related practical benefits of believing) be reduced to irrelevancy? If evidentialism is true, the critical domain theorist is left with an explanatory burden, a burden that a different kind of value-theoretic approach to epistemic justification might do better in discharging. In fact, this kind of explanatory inadequacy is seen by some as an intrinsic flaw of the critical domain approach to epistemic evaluation.[119]

The arguments to come in this chapter are aimed, in part, at problematizing epistemic exclusivism in order to exonerate the view outlined in the first three chapters with respect to this kind of explanatory worry. The evidentialist is presented with a dilemma. The horn of the dilemma relevant to a particular brand of evidentialism will

---

[119] See, e.g., the way that Grimm (2009) argues against Sosa's conception of epistemic evaluation.

depend on the evidentialist's argumentative strategy. I will argue that, if the evidentialist argues for her position in a way that utilizes a familiar kind of "no rewards principle," then, contrary to what such arguments attempt to establish, a kind of error theory about doxastic norms results. Contrarily, if the evidentialist does not appeal to a "no rewards principle," then I will argue that we either should not be convinced by the most plausible remaining principles that would exclude non-epistemic considerations from mattering in the determination of what we ought to believe or that such principles are unproblematic for the critical domain theorist. In other words, depending on the kind of argument the evidentialist gives, evidentialism is either false or does not constitute a dialectical barrier for the critical domain theorist. Either way, while the lesson of the dilemma is something less than the outright falsity of epistemic exclusivism, it absolves the critical domain theorist from any accusations of explanatory inadequacy.

Before setting up the dilemma, it will be important to make some distinctions regarding the kinds of doxastic norms in question. When, for example, proponents of SCT make "ought" claims, what kinds of claims are those supposed to be? Importantly, they are not purely descriptive; they are supposed to be *normative* in some sense or other. Suppose that I see the weather radar on my way out to the beach and I tell my friend "It ought to be sunny when we arrive." Even though I use the term "ought," I am making a descriptive claim about the weather based on the information I have. I am claiming that, given the information on the radar, it is probable that it will be sunny when we get there. The "ought" claims made in SCT are not merely descriptive in this way.

Amongst the more normative "ought" claims that one might make, there is an equally important distinction between what I will call "proper function oughts" and

"guidance-giving oughts." The former oughts are properly attributed to individuals of a kind whereby the kind has a constitutive goal or telos associated with it. Hearts ought to provide a continuous flow of blood through the body, thermometers ought to reliably produce readings in accordance with the actual temperature, teachers ought to explain things clearly to their students, etc. When a member of a kind does what it ought to in this sense, it is properly functioning vis-à-vis the relevant goal associated with the kind of thing that it is. Given the view of epistemic justification presented above, there is a sense in which it would be uncontroversial to make a connection between proper function "ought" claims and believing in accordance with one's epistemic justification. The previous chapter outlines a framework for connecting epistemic justification to a kind of proper functioning. Insofar as we restrict our evaluation to a particular propositional inquiry, someone who believes justifiably seems to be believing the way she ought to vis-à-vis inquiry, in virtue of her contribution toward the constitutive goal.[120]

The "ought" claims that at least some evidentialists make, however, are more ambitious than this. When proper function oughts apply to agents, those oughts occasionally endorse things that those agents should not do in some stronger sense. It is possible that someone ought to φ in the proper function sense when they *really* ought not φ. For example, in the proper function sense, assassins ought to stealthily eliminate their targets. However, for some particular assassin, it still might easily be the case that she

---

[120] This is one way of thinking about what many have described as the "epistemic ought," as a proper function ought explicitly tied to the epistemological critical domain. In fact, Feldman's (2001) "role oughts" seem to be a proper subset of proper function oughts. In this way, what Feldman suggests about our doxastic "ought" claims is not at all in conflict with the critical domain picture developed in these chapters.

should not kill her target at all, stealthy-ness aside. Proper function oughts are not

automatically normatively "binding" to agents in the way that other "ought" claims are.

Unlike proper function oughts, what I will call "guidance-giving oughts"

necessarily carry with them an authoritative normative force. This alternative kind of

authoritative, guidance-giving ought is notoriously difficult to describe in an informative

way but I think most readers have some kind of idea of what it consists in. David Copp

(2007) has norms like these in mind when he describes his "third grade of normativity."

Timothy Schroeder (2003) describes these kinds of recommendations as "force makers,"

Christine Korsgaard (1996) as those that "make claims on us," and Stephen Darwall

(1997) as those with "genuine deliberative weight." They are the kinds of considerations

that ultimately matter in first-personal deliberations about what to do and they are

authoritative in their guidance.[121] It is this kind of genuinely guidance-giving ought that

at least some epistemic exclusivists have in mind. Epistemic considerations purportedly

have an authority over us in their guidance about how we are to regulate our beliefs.

Another traditional distinction that should be brought out is that between

objective and subjective oughts.[122] Consider an example from Jonathan Kvanvig:

> [I]f I leave the airport knowing that there are two routes home, and resolve my
> perplexity in favor of the one that I know from experience is typically a few
> minutes faster, things still might not go well: the wreck on the chosen route might

---

[121] One may wonder, is it not the case that at least *some* proper function oughts are binding to agents in this way and that they do occasionally have an authoritative, guidance-giving force? For example, unlike assassins, we *do* think that (at least in most circumstances) teachers really ought to explain things clearly to their students and they ought to do so at least in part because of the kind of thing that teaching is. We should not deny this. The difference between being a teacher and being an assassin is that the former often has some kind of more legitimate normative standing than the latter, that there are some principles that often further generate a guidance-giving normative ought in the case of teaching and not in the case of assassinating. The point here is only that, if there is a true sense in which assassins *ought* to stealthily eliminate their targets and another true sense in which they *ought not* eliminate their targets at all, then there is some distinction in senses to be made along the above lines.

[122] For some traditional endorsements of the distinction, see Ewing (1947), Parfit (1984), and Jackson and Pargetter (1986).

> leave me stranded in traffic for three hours. While sitting there, I might begin
> kicking myself mentally: shoulda gone the other way. (2014b, pgs. 118-9)

In the example, there is a sense in which Kvanvig does exactly as he ought to have done by deciding to take the route that is usually quicker based on his experiences, but there is also another sense in which he ought to have gone the other way. The former "ought" claim attributes a subjective ought, one that relativizes to the information possessed by the agent or to some other relevant subject-specific parameter, while the latter attributes an objective ought with no such relativization.

At least some evidentialists are committed to the relevant ought being some kind of *subjective* ought. I will have more to say about this in the next section, but some evidentialists would say that the very nature of guidance-giving norms dictates that they authoritatively *guide* agents. Hence, as such, when they apply, it must be at least possible (in the appropriate sense) for them to guide the agent to which they bind, otherwise it would not make much sense to say that they are authoritative in their guidance. Because they are determined independently of any subject-relevant parameter, objective oughts are often incapable of guiding agents. Above, the fact that Kvanvig objectively ought to have gone the other way could not have possibly guided him in his deliberation about which direction to go. He had no way of knowing that there was an impending wreck waiting for him in the direction he chose. Alternatively, because of their subject-specific relativization, subjective oughts (depending on a more precise specification) often can guide us in deliberation and they have an authority in such guidance. This is not necessarily to diminish the objective ought, it is just to say that the objective ought will not be the target of discussion for much of this chapter.

I will be assuming here that guidance-giving oughts are all-things-considered normative constraints determined on the basis of guidance-giving reasons. There is much to recommend this framework. Guidance-giving reasons to φ are authoritative normative considerations that count in favor of φ-ing.[123] So as not to beg any questions for the moment, here, I take φ to range widely over anything that an agent might "do," broadly speaking; over acting, believing, desiring, or any other kind of thing that an agent "does" that can be expressed by an agent-modifying verb phrase.

In a given circumstance, one guidance-giving ought to do what one has most guidance-giving reason to do, where what one has most guidance-giving reason to do is (often, if not always) a function of the aggregation and comparison of the relevant *pro tanto* guidance-giving reasons. For example, if Kvanvig's options for going home are to take either Route A or Route B, then he ought to take Route A just in case he has most reason to take Route A. He has most reason to take Route A if the *pro tanto* reasons to take Route A outweigh the *pro tanto* reasons to take Route B in an all-things-considered aggregation and comparison. Even if he has some reason to take Route B (perhaps it is a scenic drive, among other things), those reasons might be outweighed by others including the consideration that Route A is usually quicker. Given that time is an overall weightier consideration, this would generate the verdict that he ought to take Route A. This

---

[123] Similarly to "ought," there are other, non-normative uses of the term "reason." A very common use of "reason" is to indicate a forthcoming explanation and not to indicate any kind of normative, guidance-giving consideration (e.g., "the reason the volcano erupted was the buildup of excess pressure over time"). In the final section of this chapter, I will also bring out an important ambiguity in "reason to believe." See Scanlon (1998) and Raz (1999) for some discussions of reasons to φ as normative "considerations that count in favor of" φ-ing.

connection between guidance-giving oughts and reasons indicates that they apply to the very same kinds of things, whatever those things might be.[124]

A note should be made about circularity. I have suggested that these guidance-giving norms are at least sometimes thought of as *subjective* norms, determined on the basis of some subject-specific parameter. I have also claimed that guidance-giving oughts are determined on the basis of guidance-giving reasons. Given this, one might worry that when we are determining the guidance-giving reasons one has, since those norms are sometimes subjective norms, one is bound to appeal back to the very kind of normative notion that is being determined. For an illustration of the point here, if reasons are a subjective normative constraint, then someone wishing to connect reasons and values might be attracted to the following kind of analysis: *S* has a guidance-giving reason to φ if and only if *S* has *good epistemic reason* to think that φ-ing promotes a valuable state of affairs. Comparing this kind of analysis with its more unproblematic objective counterpart (i.e., that *S* has a guidance-giving reason to φ if and only if φ-ing *actually* promotes a valuable state of affairs), it begins to look as if we are determining what reasons there are by appealing to reasons. In response, I can only give a promissory note. The final section of this chapter sets up a framework indicating that the nature of guidance-giving reasons and the nature of epistemic reasons are sufficiently distinct in a way that such worries about circularity can be avoided.

---

[124] This framework purposefully leaves the treatment of supererogation open. Saying that one ought to do what one has most reason to do is consistent with any number of ways of thinking about the relations among what one ought to do, what is permissible to do, and what would be over and above "the call of duty." For all that has been said, there may be cases in which one ought to φ even though φ-ing would be supererogatory and it would be perfectly permissible to do something other than φ. One might have most reason to φ even if one has "sufficient reason" to do something other than φ. I will say more about this in the final section of the chapter. For some further complexities that might arise in a more robust theory of rationality, see Slote (1989).

The arguments to come in this chapter are intended to cover a wide range of epistemic exclusivist views and not merely SCT above. Alternatively, an epistemic exclusivist might think that, instead of determining facts about what we *ought* to believe, epistemic considerations are exclusive in determining facts about what is *permissible* to believe. As it happens, there are all kinds of different theories about the necessary connections between epistemic considerations and doxastic norms and it would be pretty hopeless to attempt to say something informative about each one individually.[125] The common exclusivist refrain, however, is that the norms governing belief are given to us by epistemic considerations in some way or other. Here, while I will be focusing on "ought" claims in particular, claims about other kinds of connections between epistemic considerations and doxastic norms will also be relevant, especially insofar as such norms are supposed to be guidance-giving in the sense specified above. For example, in order to avoid the forthcoming dilemma, it will not do for an exclusivist to merely switch from "ought" language to, say, "permissibility" language or even "correctness" language insofar as the claims made are still supposed to be guidance-giving.

4.2: The First Horn

In this section, I discuss a particular line of reasoning that a number of evidentialists have taken up in support of epistemic exclusivism and argue that, contrary to what that line of reasoning is supposed to establish, it actually implies evidentialism's falsity. There is an explanatory burden on the epistemic exclusivist: Non-epistemic

---

[125] For some explicit claims about permissibility, see Alston (1983), Shah and Velleman (2005), and Whiting (2013). For another kind of connection between epistemic justification and what we ought to believe, see Gibbons (2013).

considerations obviously matter in the determination of the normative facts about our actions. For example, when φ ranges over actions, the *suffering* that φ-ing would cause is a consideration that matters in determining whether or not I should φ in a given situation. Given this, why think that *only* epistemic considerations matter in determining what we ought to believe? Why think that the norms governing belief and action would be so different in a way that non-epistemic considerations carry no weight at all with respect to belief? Unless we can get a good answer to that question, the truth of epistemic exclusivism would be a pretty mysterious phenomenon.[126]

In distinguishing belief from action, epistemic exclusivists occasionally appeal to differences related to what Matthew Chrisman (2008) has called the "no rewards principle" about belief (NRP): No matter how large the reward, *S* cannot simply decide to believe some proposition *p* in order to collect that reward. "[W]hen offered a reward for believing... that the U.S. is still a colony of Britain or that Hell is a bar in Chapel Hill, you cannot… just believe that the U.S. is still a colony of Britain or that Hell is a bar in Chapel Hill." (pg. 346) If we cannot decide to believe something in this way, then belief is perhaps relevantly different from at least many of the actions that we undertake. If we were offered some kind of reward to act in a certain way (say, to turn on the lights or

---

[126] One difference that has been pointed out by Kieran Setiya (2013) is that *belief* is a "static" notion (i.e., when it applies to an agent, it refers to a state that he is in, like being handsome) while *action* is a "dynamic" notion (i.e., when it applies to an agent, it is something he does or is something he can be doing in the present tense, like running). Because of this, present-tense constructions like "he is believing that *p*" seem awkward and not even well-formed. While such differences might form the basis for a more general skepticism about the kinds of norms that are applicable to belief and, while I am sympathetic to this kind of difference between belief and action, I see no reason to think that we could not relevantly shift the discussion to the more "dynamic" correlates of belief (e.g., *forming* beliefs, *maintaining* beliefs, *revising* beliefs, etc.) in order to revive the kind of question I pose above. To the proponent of epistemic exclusivism: What is so different about forming/maintaining/revising beliefs such that non-epistemic considerations do not matter at all in determining the guidance-giving normative facts of the matter?

raise our hand), we could just decide to act in that way. According to NRP, no such freedom or control exists for our beliefs.

A number of epistemic exclusivists have drawn on this purported insight. Thomas Kelly (2002), for example, claims that there is a "psychological inefficacy of practical considerations with respect to belief." (pg. 166) Kelly's way of framing NRP is particularly useful as he makes it clearer what kinds of "rewards" are supposed to be inefficacious in the formation and maintenance of belief. It is not merely the case that, say, monetary rewards would be ineffective in causing us to take up the belief that the U.S. is still a colony of Britain. The idea is that *any* perceived beneficial (but merely practical) consequence of believing would be ineffective. Even if you were offered the reward of a lifetime of happiness to take up that belief, you still could not. The inefficacious "rewards" in NRP might be any kind of perceived practical benefit of believing: a positive psychological state, a positive state of health, a positive relationship, the avoidance of harm, etc. The point is that, unlike our deliberations about turning the lights on or raising our hands, such practical benefits seemingly cannot consciously motivate us in a way that would allow us to just take the relevant beliefs up.

As Kelly points out, however, a standard pragmatist move is to attempt a denial of the difference. Take, for instance, what is recommended by Pascal (1966) in light of his famous "wager." According to Pascal, even if we cannot immediately take up a belief in God, practical considerations about the afterlife can and should motivate us to take up longer-term projects that would cause us to come to believe; we can indirectly influence our doxastic attitudes toward God's existence by choosing who we associate with, what

kinds of books we read, what kinds information we seek out, etc.[127] Moreover, as

Feldman (2000) notes, such influence and control over our beliefs need not even be

particularly long-term. Suppose that the lights are off in my house and that I am offered a

reward for believing that the lights are on. If I am really motivated by the reward and

want to believe, I can always just go and turn the lights on, causing myself to believe that

they are on! Similarly, suppose that my teenage son is out late with a group of friends

and that I am a bit worried about his safety. Seemingly, I might realize that my

uncertainty about his safety is an unpleasant doxastic state and then set about changing it

by simply calling him to ensure that he is safe.

Kelly recognizes such cases and even grants that they exhibit a relatively common

type of non-epistemic influence on belief. Yet, he thinks that there is still an important

normative difference. Unlike actions, Kelly thinks that those practical considerations that

might motivate us to change our doxastic states cannot ever be the *basis* for our beliefs in

a way that they could act as normative or justifying reasons. Put differently, while it may

be accurate in some cases to say that, in a causal or motivating sense, practical

considerations are the reason that *S* believes that *p*, it will never be the case that *S*

believes that *p* on the *basis* of those practical reasons.

Suppose, for instance, that Christie wants to believe in God for whatever practical

benefits that she thinks such a belief would bestow on her and suppose that she actually

takes Pascal's advice; she decides to hang out around believers in her free time, to

diligently read the Bible and Christian apologetics texts, go to church regularly, etc.

While Kelly is willing to grant that this aspiring Christian might end up successfully

---

[127] See Garber (2007) for a skeptical view about whether Pascal is properly read as a pragmatist.

engendering a belief in God, Christie's resultant belief would purportedly still not be *based* on those practical considerations that motivated her in the beginning. Instead, her belief would be based on the more ordinary sorts of perceived epistemic considerations that, to her, seem to speak in favor of the truth of God's existence: the arguments she learns from her readings, from the pastor, from other members of her church community, etc.

Similarly, while considerations about a reward might motivate someone to go turn on the lights in order to believe that the lights are turned on, the resultant belief that the lights are on would not be *based* on the practical considerations involving the reward. It would be based on the perceived epistemic considerations that speak in favor of the lights being on: the visual evidence that the lights are on. This is the kind of psychological inefficacy that Kelly takes to be relevant in distinguishing belief from action. As he puts it, "Although practical considerations can make a difference to what one believes, they do not do so by constituting the grounds on which beliefs are based." (pg. 174) Kelly's point entails NRP. If we could not ever believe on the *basis* of practical considerations, then, of course, we could not ever *decide* to believe anything purely on the basis of some reward that has been offered to us, no matter how large the reward is.

Nishi Shah (2006) proposes a similar principle based on a phenomenon he calls "transparency:" "the deliberative question about whether to believe that *p* inevitably gives way to the factual question whether *p*." (pg. 481) According to Shah, any time we deliberate about whether to believe some proposition, we always shift to examine the truth of that proposition and the considerations surrounding its truth.[128] For instance,

---

[128] For more on the transparency phenomenon, see Shah (2003) and Shah and Velleman (2005). Teleological versions of the so-called "aim of belief" (i.e., the claim that believers, in some way or other,

when we deliberate about whether to believe that the lights are on, we inevitably move to the question of whether the lights are on and look for considerations that speak in favor of the lights being on or not. In other words, when we deliberate about our beliefs, we automatically rely on what we perceive to be the relevant *epistemic* considerations and, as a result, practical considerations or "rewards" become irrelevant in the direct ways by which those beliefs are regulated.

Shah is careful to point out that it is in specifically conscious, deliberative contexts for which transparency holds. Shah does not mean to rule out, say, cases of wishful thinking in which a subject might, underneath the surface of consciousness, come to believe something on the basis of considerations that she would not reflectively take to be relevant to the proposition's truth (pg. 483). Transparency is merely a claim about the inefficacy of rewards and practical considerations in our *conscious* reasoning. According to Shah, we could never consciously transition to a belief in the following way: Although it is unlikely that *p*, it would be practically beneficial for me to believe that *p*; hence, I will believe that *p*. While it is a more restrictive claim than Kelly's, Shah's specification of the transparency phenomenon, again, entails NRP. If, in our conscious reasoning, we could never believe on the basis of anything other than what we perceive to be the relevant epistemic considerations, then we could never simply decide to believe something purely in order to gain a reward.

Importantly, all who appeal to NRP in defending epistemic exclusivism at least think that the kind of inability in question is a *categorical* inability, that beliefs are never deliberately based on what are merely perceived as practical considerations no matter

---

naturally aim to believe that *p* if and only if *p* is true) also sit very naturally with Shah's transparency phenomenon so long as no further normative implications are drawn.

who the believer is or what proposition in question is. There does, however, seem to be some room for disagreement about the relevant modality involved in the inability claim. For example, Shah takes the transparency phenomenon to hold out of conceptual necessity (or, at the very least, he thinks that certain conceptually necessary truths explain the relevant inability; see Shah 2003, pgs. 465-74).[129] This strong modal claim is not necessarily required for the kind of argument he and others wish to give in favor of epistemic exclusivism.[130] Kelly, for example, is more agnostic about the modality involved but he still attempts to argue for exclusivism on the basis of NRP. For the purposes here, the modal force of the inability claim, whether it be conceptual, metaphysical, nomological, or something else, is not important so long as the inability is at least as modally strong as the kind of practical inability I specify below.

Although there are a number of ways of framing an argument from NRP to epistemic exclusivism, I will be focusing specifically on Shah's (2006) version here as I think it is the clearest and most explicit form available. Inspired by Bernard Williams (1981), Shah argues for what he calls "the deliberative constraint" on reasons: R is a reason for X to φ only if R is capable of disposing X to φ in the way characteristic of R's functioning as a premise in deliberation whether to φ. (pg. 485) The idea is this: In order for some consideration to be a normative reason for someone to act or believe in a certain way, that consideration must be capable of figuring into one's deliberation *as a reason* about how to act or believe. In order for some consideration to be a reason for me, it is a

---

[129] See Adler (2002) and Adler and Hicks (2013) for some agreement about the conceptual necessity of the inability.
[130] See Shah (2006, pg. 484).

necessary condition that there is not some unalterable feature of my psychology that would prevent me from deliberately reasoning on the basis of it.

When this deliberative constraint on reasons is applied specifically to beliefs and is combined with the transparency phenomenon (or, for our purposes, NRP), the result is that practical considerations can never be normative reasons to believe.[131] If, for all *S*, a consideration must be able to figure in *S*'s deliberation about whether to believe that *p* in order for it to be a normative reason for *S* to believe that *p* and if, for all *p*, purely practical considerations cannot ever figure our deliberations about whether to believe that *p*, then the result is that *S* cannot ever have a purely practical reason to believe that *p*.

Shah attempts to support the deliberative constraint on reasons from two directions. First, he thinks that the "best explanation" for the phenomenon of transparency entails the deliberative constraint (pgs. 488-9). According to Shah, transparency is best explained by the fact that the concept of belief comes with a constitutive norm of correctness whereby a belief is correct just in case it is true. Purportedly, insofar as one engages in doxastic deliberation and explicitly draws on the concept of belief, one is already committed to believing some content if and only if that content is true, explaining why only considerations related to the truth of the content can move us to believe.

This explanation is thought to entail the deliberative constraint on reasons. Shah puts it in the following way:

> My explanation of transparency conceives of deliberation as a norm governed activity. Deliberation is reasoning aimed at issuing in some result in accordance with norms for results of that kind. Deliberating about whether to φ is reasoning aimed at issuing or not issuing in φ-ing, in accordance with norms for φ-ing. The

---

[131] See Rescher (1984), Foley (1987), Nozick (1993), and others for explicit endorsements of the view that there *are* practical reasons to believe.

deliberative constraint on reasons naturally flows from this conception of deliberation. Deliberating whether to φ is not merely one amongst other causal routes to securing results of kind φ; it is the most explicit route by which we are guided by the application of norms in reaching results of kind φ. Reasons for φ-ing are considerations which indicate whether φ-ing would be correct according to the norms for φ-ing. Deliberation whether to φ thus is precisely the place to look if we want to understand what can be a reason for φ-ing, and not just a mere cause of φ-ing. (pg. 489)

Now, independently of whether this purported entailment relation holds between Shah's "best explanation" of transparency and the deliberative constraint on reasons, this line can only support the deliberative constraint if that "best explanation" of transparency that Shah provides us is, indeed, a genuine explanation of transparency. Asbjørn Steglich-Petersen (2006) gives us good reasons to doubt that it is a genuine explanation. Instead of the concept of believing, consider the concept of promising. If it is plausible to think that the concept of believing contains normative constraints, then it is equally plausible to think that the concept of promising does so too. For the concept of promising, one might think that it contains the constraint that it is appropriate for one to make a promise just in case one intends to carry out that promise. However, even if the concept carries with it a constraint like this, it is still quite possible for us to simply ignore that constraint in our deliberations about whether or not to make some particular promise. Even if the constraint is part of the concept of promising, we are not necessarily committed to deliberately regulating our promises in ways that would conform to it. It is quite possible for us to deliberate and to decide to make a promise fully knowing that we do not intend to carry it out. Hence, it seems that even if the concept of believing carries with it the normative constraint that Shah imagines, it still would not explain transparency.

Consider another concept that is perhaps even more compelling as a counterexample to Shah's explanation: the concept of assertion. Again, it might be plausible to think that the concept of assertion contains a normative standard of correctness; as a matter of conceptual necessity, an assertion is correct just in case the content of the assertion is true. However, when we deliberate about whether or not to make an assertion at a given time, we might easily ignore that normative constraint and it is quite possible for us to deliberately make an assertion for reasons independent of its truth.[132] If it is possible for us to do so in our deliberations about our assertions, then why is it impossible for us to do so in our deliberations about our beliefs? One might be tempted to appeal to the phenomenon of doxastic transparency but, of course, that is exactly the phenomenon that Shah is trying to explain in the first place.

Independently of his "best explanation" entailment, Shah attempts to provide another line of support (a better line of support in my estimation) for the deliberative constraint on reasons (see, e.g., pg. 486). As he sees things, the kind of normative reasons in question are *guidance-giving reasons* in the way described above (see pgs. 156-8 in this chapter), as authoritatively figuring into what we ought to do or believe in the guidance-giving sense. According to Shah, if a kind of consideration could not possibly guide us in deliberation, then that kind of consideration automatically departs from the nature of guidance-giving normative constraints. The nature of normative reasons dictates that they are guiding and that they are authoritative in their guidance. Clayton Littlejohn (2011, pg. 122) agrees with Shah: "nothing can be a reason unless it can figure in reasoning." More specifically, nothing can be a guidance-giving reason

---

[132] This example is related to a similar point made by Stephanie Leary (2016) in response to the so-called "wrong kinds of reason" problem. I will discuss her point in detail later in this chapter.

unless it could actually guide us in our deliberation about what to do or believe.[133] This constraint might appropriately be thought of as the guidance-giving reasons analogue of the "ought implies can" principle. If it is true that S (guidance-giving) ought to φ, then it is possible for S to φ. Similarly, if some consideration R is a (guidance-giving) reason for S to φ, then it is possible for S to deliberate on the basis of R in determining whether or not to φ. If that is right, then Shah would have some support for the deliberative constraint on reasons. When combined with transparency, there can be no practical reasons to believe.

My goal in the remainder of this section is not actually to criticize this line of argumentation. I find it plausible that, if NRP and the deliberative constraint are both legitimate, then there are no practical reasons to believe anything in the guidance giving sense. However, I do want to criticize a further implication that has been drawn. Shah, for example, does not merely take his argument to rule out practical reasons. He goes further; because practical reasons have been eliminated, he takes such an argument to support a guidance-giving form of evidentialism, the view that epistemic considerations are exclusive in determining the guidance-giving norms that govern our beliefs. Here, I will offer a conflicting account: *If* NRP is true and the deliberative constraint on reasons is legitimate in virtue of the guidance-giving nature of the norms in question, then we

---

[133] The possibility of *rational dilemmas* puts some pressure on this idea. For instance, if it turned out to be possible for S to have most guidance-giving reason to *both* φ and not-φ (or even for S to have most reason to both φ and ψ where ψ-ing entails not-φ-ing), then it would seemingly be possible for S to have guidance-giving reasons to do the impossible, making it impossible for such reasons to actually guide S in deliberation. At the very least, it seems as though Shah is comfortable ruling out impossible states of affairs as genuine options to be deliberated about. In the way Shah is thinking about them, all-things-considered, guidance-giving normative constraints should be *univocal* and should not demand contradictory states of affairs. See Conee (1982) and Zimmerman (1996) for some agreement here. Shah seemingly follows Thomas Hill Jr. (1996) in thinking that, if a guidance-giving normative theory can eliminate "gaps" in guidance (of which the demand for contradictory states of affairs is an instance), then it is theoretically desirable to do so.

should be error theorists about doxastic norms in a way that rules out evidentialism just as much as it rules out pragmatism.

Shah thinks that the guiding-giving nature of the norms in question supports the deliberative constraint on reasons. If we are to accept this, then it seems as though we should also accept what would appropriately be labelled the "practical possibilities requirement" on guidance-giving norms (PPR): Guidance-giving norms (oughts, reasons, etc.) apply to S's φ-ing at time t only if, for S, there is at least one practically possible alternative to φ-ing at t. Applied specifically to guidance-giving oughts, S ought to (or ought not to) φ at t only if, for S, there is at least one practically possible alternative to φ-ing at t. For guidance-giving reasons: S has (at least some) reason to φ (or not to φ) at t only if, for S, there is at least one practically possible alternative to φ-ing at t. In a given moment, if there are no practically possible ways for us to do anything other than what we *actually* end up doing, it would be a mistake to think that any real guidance could be given to us. What we actually do may, of course, be better or worse than some imagined alternative, but unless that alternative is a real practical possibility, the betterness or worseness of it would not be able to give us any real direction or guidance. If there are no practically possible alternatives to what we do, our deliberations about what we ought to do would be based on a mistake; trying to answer the question of what, among the practical possibilities, we ought to do or have most reason to do would involve a faulty assumption. This is not say that it is always *worthless* to deliberate about impossible states of affairs (practically or otherwise), it is only to say that, if the goal is for such deliberation to issue in guidance-giving norms, then there is something defective about it.

Hence, if there are no practically possible alternatives to what we do, then guidance-giving norms are simply inapplicable to us.

PPR is more than a mere "ought implies can" principle.[134] PPR is more appropriately thought of as representing an "ought implies can-not" kind of principle: *S* ought to φ only if it is practically possible for *S not* to φ.[135] It will be worthwhile here to further specify what it means for something to be "practically possible" in this sense. These practically possible alternatives are a set of states of affairs that an agent has a relevant kind of control over. Consider an example: Suppose that I am running a bit late for an important meeting. If someone were to say that I ought to hurry up (in the guidance-giving sense), then there is an implication that it is somehow within my control whether I hurry or not. It is implied that I could decide whether I hurry or not, that I have a specific kind of power and ability at that time with respect to those alternative states of affairs or courses of action. Generally speaking, the states of affairs over which I have this relevant kind of control are the ones that comprise the practical possibilities in a given moment.[136]

Hence, PPR is the idea that, in order for guidance-giving norms to apply to φ-ing, we need some kind of relevant control over whether we φ or not. What, though, is this "relevant" kind of control? Taking a hint from Alston (1988b) and others, we can make some distinctions between types of control. For example, I might have a kind of basic

---

[134] For some general discussion about the "ought implies can" principle, see Montefiore (1958) and Howard-Snyder (2006).

[135] There are some striking similarities between PPR and what is proposed in Wedgwood (2013b). Guidance-giving oughts somewhat resemble Wedgwood's "non-trivial" oughts.

[136] While I occasionally use the language of what I can "decide" to do in order to describe practically possible alternatives, it should be noted that I make no assumptions about what "free decisions" or about what "free will" consists in. The claims below will apply independently of the particular theory about free will that turns out to be correct, at least among the most plausible candidates. See Pereboom (2009) for a sample of the possibilities here.

and immediate control over whether I φ or not. I have a basic and immediate control over my φ-ing when I can "just φ," as it were, in virtue of the kind of direct and temporally immediate power I have over my body and limbs; e.g., I can just raise my hand, turn my head, kick my legs, blink my eyes, etc. I might also have a non-basic but still immediate control over whether I φ. For instance, if my hand is on the doorknob of my front door, I have a non-basic but still immediate control over whether I open the door. The control is immediate because I can open the door in that very moment, but it is non-basic because my opening the door would obtain in virtue of the control I have over a more basic action type (e.g., turning my wrist in a clockwise manner).[137]

It seems as though both of these kinds of immediate control (basic and non-basic) are capable of generating the kinds of practically possible alternatives that are necessary for the application of guidance-giving norms. Suppose that I am at my front door with my hand on the doorknob, that I am expecting a package or visitors and I hear the doorbell ring. I have a relevant kind of control over what I do in a way that guidance-giving norms can apply. For example, I have control over whether I open the door or not and, given that there are no defeating considerations, it would be appropriate to say that I have good reason to open the door and even that I ought to open the door. The same holds for the more basic action type of turning my wrist in a clockwise manner. I have reason to do so and ought to do so. Guidance-giving norms can apply to what I do because there is a relevant practical possibility space generated by the kind of control that I have.

---

[137] See Danto (1965) for the original distinction between basic and non-basic action types.

Orthogonal to the basic/non-basic distinction, I might also have control over what I do in a way that is more or less long-term. For instance, when I hear my doorbell ring, I still have a kind of control over whether I open the door or not even if I am currently sitting on my living room couch. If I am expecting a package or visitors, we would still not hesitate to say that I have good reasons to open the door and that I ought to open the door even if I cannot simply open the door immediately in that exact moment. There are certainly gradations in how long-term a token of control might be, but these gradational differences really do not seem to be relevant to the generation of the necessary practically possible alternatives. It does not really seem to matter, for example, where I am relative to the door or how long it would take me to get to the door in order for me to have the relevant control over my door opening. In this particular example, what really seems to matter in generating the practical possibility space is that I *could* put myself into a position such that I would have an immediate control over whether I open the door or not. *If* I were to do certain things and carry out certain actions that I do have a more immediate control over (e.g., getting up from the couch, walking toward the door, etc.), I would eventually be in a position such that I would have an immediate control over whether I open the door or not. Generally speaking, this kind of long-term counterfactual control over φ-ing, the fact that an agent *would* have an immediate control over φ-ing at a later time were she to exercise her relevant control over certain other things, is enough to generate the necessary practical possibility space for guidance-giving norms to apply to φ-ing.[138]

---

[138] Similarly, other kinds of counterfactual considerations can also show us why *consciously intending* to φ or not or *consciously deliberating* about whether to φ is not necessary for the application of guidance-giving norms. For example, even before I actually turn the ignition on in my car, it would be appropriate to say that I have a relevant kind control over that action even if I end up somewhat mindlessly going through

Long-term control should be contrasted, then, with mere long-term probabilistic influence. For instance, there may be a sense in which I have *some* control over whether I win the lottery or not. At the very least, I have an immediate control and long-term control over certain things that would increase my chances of winning. I could get up from my chair, walk out of the building, get in my car, drive to the nearest lotto vendor, and purchase as many tickets as I please. Yet, despite the fact that I have control over these things that would increase my chances, at no point would I ever have an immediate control over whether I win the lottery or not (unlike the kind of long-term control I have over my door opening).

This is a relevant distinction because it seems to carve a line between the kind of control that generates the relevant practical possibility space and the kind that does not. Mere long-term probabilistic influence is not a strong enough form of control and, hence, guidance-giving norms do not apply directly to what we merely have long-term probabilistic influence over. This helps explain the awkwardness of the claim that someone *ought* to win the lottery (in the guidance-giving sense). It very well may be *good* for me to win the lottery, but that goodness would only seem to create direct guidance-giving normative constraints on the things that I have a more robust form of control over. If it would be all-things-considered good for me to win the lottery, perhaps I then ought to get up and go to the store to buy some lottery tickets, but it would be

---

the motions of getting to work and I do not consciously think about turning on my ignition at all. A relevant kind of control is exhibited, not because of any kind of conscious intention is formed, but because of the action's *counterfactual* relations to my more conscious deliberations: *If* I were to consciously think about turning my ignition on or consciously intend to turn my ignition on, I could put myself in a position such that I would have an immediate control over doing so. Again, this is enough to generate the relevant practical possibility space in order for guidance-giving norms to apply. Hence, the kind of general worry about doxastic intentions brought out by Buckareff (2006) is orthogonal to the kind of control utilized in PPR. See Steup (2008) for some agreement here.

strange to further claim that I also ought to win the lottery. I just do not have the relevant kind of control over whether I win the lottery or not and, hence, such guidance-giving norms do not directly apply to my winning the lottery. Telling me that I ought to win the lottery would be to make a category mistake just as much as it would be to tell an Army veteran with a debilitating leg injury that they ought to regain the use of their leg. While it might make sense to say that the veteran ought to take care of herself and go to physical therapy and the like, the lack of control over the efficacy of such rehabilitation rules out the direct application of guidance-giving norms.

PPR, then, imposes a constraint on the application of guidance-giving norms like oughts and reasons. This constraint conflicts with other possible ways of thinking about these norms. Mark Schroeder (forthcoming), for example, proposes an account about the "unity of reasons" wherein reasons might apply to anything at all that an agent might "do," anything at all that can be expressed by a modifying verb phrase. According to a view like this, since an agent might hurry up, open a door, believe that *p*, have a mental breakdown, have a reflex reaction, spontaneously combust, win the lottery, or "do" any other number of things, an agent might thereby have reasons to (or reasons not to) hurry up, open a door, believe that *p*, have a mental breakdown, have a reflex reaction, spontaneously combust, win the lottery, etc. The point of PPR is that, insofar as reasons are the kinds of things that aggregate in generating guidance-giving norms like oughts, they will not apply to anything and everything that we are capable of "doing," but only to those things that we have a relevant kind of control over. Not all things that we can "do" will be subject to the direct application of guidance-giving oughts and reasons.

With this groundwork laid out, the relevant argument for this section can finally be given. If Shah and other evidentialists are right about NRP, then, contrary to what such evidentialists try to establish, evidentialism would be false. If both NRP and the considerations underlying the deliberative constraint on reasons are legitimate, then epistemic considerations cannot be exclusive in determining the guidance-giving normative facts about belief because there would be no such facts. Guidance-giving norms would simply be inapplicable to believing in the same way that those norms are inapplicable to all of the other things that we "do" that we do not have a relevant kind of control over. Shah's claims that, in deliberation, we must always believe on the basis of what we perceive to be the relevant epistemic considerations and that we are incapable of basing our beliefs on any other kind of consideration. If that is right, then, as I will argue, we do not have the relevant kind of control over how we base our beliefs in a way that is necessary for the application of the kinds of norms Shah wishes to apply (i.e., guidance-giving norms). For belief, there would be no relevant practical possibility space, no relevant alternatives through which we might be guided.

Consider the example involving Christie, the aspiring theist. Since she thinks that believing in God would benefit her in certain ways, she desires to believe in God and sets about doing so. She decides to read the Bible regularly, go to church, make friends with other churchgoers, etc. and eventually, let us suppose, that she is successful and that she becomes a believer. If Shah and other evidentialists are correct about NRP, even though she ends up successfully believing (at least partly) because of her efforts, at no point in that process does she ever have a relevant control over whether she believes or not. When she initially comes to think that a belief in God would be beneficial, she does not

have control over her doxastic state because, at that moment, she is completely at the mercy of however the relevant epistemic support relations happen to strike her. How she perceives those epistemic support relations at that moment is just something that is not up to her. Epistemic support relations striking us in certain ways are *passive phenomenological experiences*. If those passive experiences determine our doxastic attitudes in a given deliberative moment, then we cannot have a relevant control over our beliefs.[139]

It is exactly because of this that Christie purportedly (according to proponents of NRP) cannot believe in God at the outset and it is why she initiates a longer-term project of putting herself in the right situation to believe. But, even when she carries out that long-term plan, again, if Shah and others are correct, she is still always going to be at the mercy of however she perceives the epistemic support relations and whatever she happens to find truth-indicative in a given moment. Even when she becomes convinced by the arguments that she learns, her doxastic state is still not something that is within her control. Even though she happens to be convinced, the arguments she learns might just as easily strike her differently and, if Shah and others are right, this would automatically inhibit belief. It is just not up to us, in a given moment, whether we perceive some proposition to be epistemically supported or not.

Christie's long-term plan to believe that God exists is far more like a plan to win the lottery than it is like a plan to, say, open her front door while she is sitting on the couch. Even though she would need to carry out multiple actions before actually doing so, she still has a relevant kind of control over whether she opens her front door or not.

---

[139] Another way of putting the point is that, if NRP is true, then belief, even in the form of judgment, is not an "exercise of agency" as Sosa (2015) would have it.

Hence, guidance-giving norms can apply to her door opening. Alternatively, she does not have this kind of control over whether she wins the lottery or not. She can influence her chances of winning the lottery by doing other things that she does have control over, but she is never going to be in a position to directly control whether she wins or not. Because of this, guidance-giving norms cannot apply directly to her winning the lottery. While it might be the case that she ought to (or ought not) go out, get in her car, and buy some tickets, it would be a category mistake to say that she ought to (or ought not) win the lottery, at least in the guidance-giving sense.

Something similar holds for her long-term plan to become a theist. While it might be appropriate to say that she ought to (or ought not) read the Bible, go to church regularly, and make friends with other churchgoers, it would be a mistake to directly say that she ought to (or ought not) believe that God exists. If Shah and others are right about NRP, there are no guidance-giving normative properties that attach directly to believing particular contents.

What about the following kind of case?: Suppose that Adam consciously but hastily forms the belief that p on the basis of an insufficient set of epistemic considerations while those considerations actually support the truth of $\neg p$. Even though he is wrong, if NRP is true, Adam must still have *thought* that those insufficient considerations support the truth of p in order for him to base his belief on them. However, was it not possible for him to, say, take his time, look at the epistemic considerations more carefully, and perhaps see that they support the truth of $\neg p$? By NRP, this might have caused him to believe that $\neg p$. Would that not indicate that there is

a practical possibility space involved, one that makes it appropriate to say that Adam ought to believe that $\neg p$?

While there may be a kind of practical possibility space involved, it is not a *relevant* practical possibility space, one that indicates that Adam has a requisite kind of control over his belief. The case is again analogous to an "ought" claim about winning the lottery. For instance, even if I do not currently have any lottery tickets, it is still practically possible, in at least some sense, that I win the lottery. I could, after all, go out and buy some tickets which might make me a winner. The point here, however, is that, even if I do go out, do my due diligence, and buy as many lottery tickets as I can, it is still not completely up to me whether I win. I am at the mercy of however the lotto balls are picked. I may win. Then again, I may not.

Similarly, even if Adam were to do his due diligence, sit down and contemplate the epistemic support relations more than he actually does, it is still not completely up to him whether he will end up believing that $\neg p$ (at least if NRP is true). He would still be at the mercy of however those epistemic support relations happen to strike him in that moment. They may strike him as supporting the truth of $\neg p$. Then again, they may not. That is simply not going to be up to Adam in the relevant way.

Consider a different kind of case: Suppose that I believe that p and that I consciously apprehend an epistemic support relation between p and q, namely that p implies q. Is it not the case that there is a practical possibility space open to me? Even consistent with NRP and the apprehended support relation, it seems as though I still might do one of two things: I might (1) simply take up the belief that q or (2) reject q and also give up my belief that p. Are these not different options in a way that makes room

for the application of guidance-giving oughts? When we apprehend support relations in this way, do we not often have a choice about whether we accept the conclusion or reject the relevant background assumptions?[140]

Again, if NRP is true, even though it would be appropriate to say that there are "options" here, there are not options in a way that indicates a relevant kind of control. If I have the belief that p and apprehend that p implies q, whether I simply believe that q or give up my belief that p depends on features of my situation that are just not up to me. If NRP is true, then whether I believe that p depends on whether or not I think p is epistemically supported and whether I believe that q depends on whether or not I think that q is epistemically supported. Per NRP, prior to my apprehending that p implies q, I must think that the epistemic considerations support the truth of p since I already believe that p. Now, when I come to apprehend the entailment relation, whether I end up believing that q or giving up my belief that p will depend on whether the perceived epistemic support relations surrounding p significantly change, and this is not something that would be within my control. When I apprehend the entailment, if I perceive q to be a rather epistemically innocuous proposition (which is not something that is within my control), then I will perceive q to have just as much epistemic support as p, engendering the belief that q. Contrarily, if I perceive q to be an epistemically strange or even absurd proposition (which is also something that is not within my control), then the perceived epistemic support for p would change in virtue of apprehending the entailment relation, perhaps so much that I no longer think that p is sufficiently supported, purportedly causing me to give up my belief that p. Hence, when I recognize that p implies q, if NRP

---

[140] Thanks to David DiDomenico for pushing me on this case.

is true, while there are different ways that my deliberation *could* turn out in some sense, these are not possibilities that I have a relevant kind of control over.

The point being made here is not an entirely new one, though I do not think that it has been brought out so explicitly in the evidentialism/pragmatism debate. Alston (1988b), Plantinga (1993b), and many others, for example, discuss the normative implications following from deontological views about epistemic justification (which seemingly do involve guidance-giving norms) and they do discuss what such views seem to assume about voluntary control. However, independently of deontological views about epistemic justification, the point here is that NRP seems to commit us to a kind of involuntarism that rules out the possibility of a distinctively *guidance-giving* form of evidentialism or epistemic exclusivism (or, more generally, *any* kind of directly guidance-giving doxastic normativity).

Jonathan Adler, a self-proclaimed evidentialist, also seems to recognize this implication: "If, under ideal conditions, I cannot help but believe (not believe) when I recognize that the evidence establishes (fails to establish) that *p*, it makes no strict sense to say that I ought… to believe *p*." (2002, pg. 51) Since Adler is also friendly to the kind of NRP expressed in the antecedent, he moves to an evidentialist view under which it is merely "proper" or "in accord with the concept of belief" for one to believe in accordance with the evidence. Because this alternative version of evidentialism does not make an "ought" claim and because it purportedly no longer contains any "directive to action," Adler thinks that it is no longer subject to any worries about the lack of control. While a move like this might make his evidentialism consistent with both NRP and PPR, I think we should still be a bit worried about it. Is the "proper"-ness of believing in accordance

with the evidence really not supposed to be guidance-giving? On one hand, if it is still supposed to be a guidance-giving norm, then Adler's alternative formulation of evidentialism would still not have avoided the worry brought out here (see pg. 162 above). On the other hand, if it is not a guidance-giving norm, then it is unclear why we should care *exclusively* about it when we are actually deliberating about how to regulate our doxastic states, either directly or indirectly.

Now, if there were even *some* circumstances in which we have a requisite control over how some of our beliefs are formed or maintained, then guidance-giving norms might apply directly to those token doxastic attitudes. However, this is exactly what proponents of NRP will deny. NRP, again, is supposed to act as a categorical inability. Purportedly, we cannot decide to base our beliefs on anything other than the epistemic considerations at hand and, as I have pointed out here, this would mean that we never have the necessary control over our beliefs for guidance-giving norms to apply directly to them. Hence, if it is the case that epistemic considerations are exclusive in determining the normative facts about belief, then it cannot be for the reasons that Shah and others give to us by appealing to NRP.

As a final note, nothing that has been said above seems to imply anything in particular about normative *responsibility* and its direct application to our doxastic attitudes. It might be consistent with the above to claim that, even if NRP is true, we are still normatively *responsible* for the beliefs we hold, even if guidance-giving oughts and reasons do not apply to them. Because the notion of responsibility serves its own unique normative-evaluative function and it is tied in the way it is to praiseworthiness and

blameworthiness, it may turn out to be the case that we can be responsible for things that we currently only have a probabilistic influence over.

Take an example: Suppose that I am currently swimming in a tank full of sharks. It might be plausible to think that I would be normatively responsible and blameworthy for my being eaten by sharks even if I cannot directly control, in that given moment, whether I am eaten by sharks or not, especially if I had an antecedent control over whether I jumped into the tank or not. After I am in the shark tank, even though I merely have a probabilistic influence over whether I am eaten or not (surely I could do certain things to better my chances), and even though it would be strange to say that I should (in the guidance-giving sense) refrain from being eaten by sharks, I might still be normatively responsible or blameworthy for being eaten by sharks because I could have prevented my own bad situation by performing some action that I previously did have control over (i.e., refraining from jumping in the shark tank).

Similarly, perhaps it will be plausible to think that someone is still normatively responsible for their beliefs even if it would be a category mistake to make any direct, guidance-giving claims about those beliefs. Could an NRP-friendly version of epistemic exclusivism be revived by shifting to claims about our doxastic responsibilities? Could it be, for example, that we are always normatively *responsible* for not taking on beliefs that are epistemically unjustified? Perhaps, but a claim like this would take independent support over and above what has been provided by Shah and others who use NRP as a basis for epistemic exclusivism. *Why* should we think that we are always normatively responsible for making our beliefs fall in line with what we are epistemically justified in believing? Why rule out other kinds of considerations? Are other, non-epistemic kinds

of considerations automatically ruled out because they are instances of the so-called "wrong kind" of reasons? Are other considerations ruled out because they never promote anything of genuine value? Are the facts about our doxastic responsibilities just primitive normative facts?

In the next section, I will survey what I think are some of the most plausible alternative lines toward evidentialism. While it will be impossible to cover every possible kind of argument that an evidentialist might give, I think those who would hope to argue for a kind of evidentialism about our doxastic responsibilities are likely to need arguments similar to the ones considered there and, hence, the discussion will apply, *mutatis mutandis*, to such a position.

4.3: The Second Horn

The lesson of the previous section is that, if NRP is a categorical truth about how we form and maintain our beliefs and that is used in tandem with a deliberative constraint on reasons, then any guidance-giving form of evidentialism must be false because we would never have the required kind of control over whether we believe something or not in order for those norms to apply directly to our doxastic states. What, though, if we do not rely on the truth of NRP? What if we leave it as an open possibility that, at least in some cases, we do have the required form of control over our beliefs? Are there any other kinds of argumentative strategies that might still push us to be exclusivists about epistemic considerations? Though it would be impossible to survey all of the possible arguments that an evidentialist might give, here I will discuss what I think are, at least on the surface, the most promising basic lines independent of a reliance on NRP. I end up

either rejecting their dialectical efficacy or I argue that they are unproblematic for the critical domain theorist and, hence, the critical domain theorist is exonerated from related explanatory worries (see pgs. 155-6 above).

If we are not going to be committed to a categorical form of NRP, it would be worthwhile to bring out some cases in which there may be a relevant kind of control over belief. Consider Christie, once again. Suppose that Christie realizes that the purported evidence for God's existence is questionable and that she currently lacks sufficient epistemic support for believing that God exists. Yet, Christie still comes to think that believing in God would be good for her. Proponents of NRP claim that Christie is thereby incapable of basing a belief in God on such practical considerations and, if she wants to believe, she must undertake some kind of long-term project that would better her chances of perceiving the epistemic support relations differently. But, is that really true? Would it really be impossible for Christie or for anyone else to, say, find some of the stories in a relevant religious text to be inspiring, to recognize that her family and friends believe in the existence of God and that many of her relationships would be harmed if she could not come to believe, and then consciously base a belief in God on such considerations? At the very least, it is not completely obvious that this is a psychological impossibility. If it is a real, psychological possibility for her to believe in this way, then it seems as though Christie might have a direct kind of control over what doxastic state she takes up with respect to the existence of God; she could either (1) suspend belief with respect to the existence of God in accordance with the perceived epistemic considerations or (2) believe in the existence of God in accordance with the perceived practical considerations.

Anecdotally, I know of a number of theists who at least *think* that they believe in the existence of God for practical reasons very similar to these. I am reminded of a particular conversation with a theist who told me that his belief in God is "not a matter of logic," openly claiming that he lacks epistemic justification for it. Nonetheless, he purportedly holds on to his belief because of the kind of life he wants for himself, for his family, and for other people around him. Now, *perhaps* people like him are deluded about their own reasons for believing and that they are making a mistake in thinking that they are believing for practical reasons related to their own social situations. *Perhaps* people like him do not really believe at all and are only pretending to believe for the practical benefits.[141] However, if we take their conscious reflections at face value, cases like these seem to be a kind of counterexample to a categorical form of NRP. In certain forms, fideism is inconsistent with NRP.

Consider another kind of case: Suppose that Anne is trying to decide whether to believe a rumor that she heard about a good friend of hers, a rumor about her friend being guilty of some crime. Also suppose that Anne's friend denies the charge. In a case like this, it is not obvious that Anne would be incapable of consciously coming to believe that her friend is innocent for what she perceives to be non-epistemic reasons related to the nature of their friendship (e.g., because she thinks that good friends ought to trust one another or because she thinks that her friend *deserves* her trust). Even if Anne perceives the overall epistemic considerations to be murky, pointing neither towards guilt nor innocence, it seems as though she might have a kind of direct control over what she believes. She could either (1) suspend belief about whether her friend is guilty or

---

[141] Knowing this person's character, I am very skeptical of this interpretation, at least in his case.

innocent in accordance with the perceived epistemic considerations or (2) believe that her friend is innocent in accordance with the perceived non-epistemic considerations regarding their friendship.[142]

The point here is not to make any particular claim about whether or not we actually have control over any or all of our beliefs. The point is the following: If the evidentialist is not going to rely on the categorical truth of NRP, then she will need some further argument in order to support her view that, even in cases like those described above where there is at least a *prima facie* kind of direct control and there are seemingly strong non-epistemic reasons to believe contrary to what the epistemic considerations would support, we still always ought to take up those doxastic states that are in accordance with the relevant epistemic considerations. The evidentialist needs some further argument in order to convince us that Christie, Anne, or anyone else should always form and maintain their doxastic states in accordance with the evidence, broadly speaking.[143]

In such cases, if we are not going to rely on NRP, why should we still think that the epistemic considerations are exclusive? Some evidentialists have thought that, even if there *are* practical reasons to believe, those reasons are defective in some sense or other

---

[142] See Keller (2004) and Stroud (2006) for more on friendship and belief. It should be noted that claims about control in these kinds of cases would be more ambitious than the merely metaphysically possible creatures described by Bennett (1990). Bennett's Credamites are imagined to have control over their beliefs, but only because they are capable of completely forgetting the practical basis for their beliefs immediately after those beliefs have been taken up. Here, Christie and Anne are imagined to be in real, practically possible scenarios where they might not only take up a belief for practical reasons, but do so with a subsequent full awareness that such considerations are not truth-indicative. See Ginet (2001), Frankish (2007), and Weatherson (2008) for claims that we do, in fact, at least sometimes have a relevant control over our doxastic states. Also see McCormick (2015).

[143] There are many other kinds of cases that have been brought out as candidates for thinking that we sometimes have good reasons to believe contrary to what the epistemic considerations support. See, e.g., the kind of positivity bias recommended in Preston-Roedder's (2013) "faith in humanity." These kinds of cases, among a number of others, are discussed at length in Hazlett (2013).

and, as such, they do not carry with them the kind of normative force that is supposed to come with good epistemic reasons. There are the "right kind" of reasons and the "wrong kind" of reasons. In this direction, a number of authors have utilized a kind of ambiguity when we speak of "reasons to believe" some particular content. Christian Piller (2006) describes the relevant ambiguity as a difference between "content-related" and "attitude-related" reasons, Derek Parfit (2011) as a difference between "object-given" and "state-given" reasons, Pamela Hieronymi (2005) as a difference between "constitutive" and "extrinsic" reasons, Jonathan Adler and Michael Hicks (2013) as a difference between "reasons to believe" and "reasons for believing," and so on.

Because I am partial to Piller's terminology, I will describe the distinction here as a difference between "content-related" and "attitude-related" reasons and will bring out the subsequent argument for epistemic exclusivism in those terms as well. The distinction is this: For attitudes and states with constitutive "correctness" or "fittingness" conditions, content-related reasons, appropriately, are considerations directed at the *content* of the attitude or state with respect to its correctness or fittingness while attitude-related reasons, appropriately, are those directed at the attitude or state itself, independently of the correctness or fittingness of the content. Consider belief in particular: By many accounts, token beliefs are "correct" or "fitting" when the contents of those beliefs are true. Hence, content-related reasons for believing that *p* would be considerations that count in favor of the truth of *p*. On one hand, insofar as epistemic reasons are supposed to count in favor of the truth of a target proposition (either by probabilifying it or providing some other kind of support), they are content-related. On the other hand, insofar as they exist at all, merely practical reasons to believe would be

attitude-related. In a slogan, epistemic reasons to believe that *p* are considerations that count in favor of the truth of *p* while other, non-epistemic reasons to believe that *p* are considerations that count in favor of believing that *p*, independently of the truth of *p*.

Now, some who point out this distinction use it in support of epistemic exclusivism. For instance, Hieronymi (2005, pg. 448) argues that merely attitude-related reasons "are not 'really' reasons for believing *p*" while Skorupski (2010), Parfit (2011), and Whiting (2014) all independently argue that merely attitude-related reasons are not "genuinely normative" in some sense or other. If that were correct, then exclusivists would have a principled way of ruling out non-epistemic reasons in determining the normative facts about our beliefs.

Why might someone think that attitude-related reasons are not genuinely normative? Consider an example from Stephanie Leary (2016, pg. 531): A state of admiration is "correct" or "fitting" if and only if the object of admiration is genuinely admirable. Because of this, it seems to us that considerations independent of whether something is admirable do not constitute normative reasons to admire it. For instance, so the argument goes, the fact that admiring my mother would increase my inheritance is not a normative reason to admire her. Only considerations relevant to my mother's character and her admirability could give me genuinely normative reasons to admire her. *Mutatis mutandis*, since a belief that *p* is "correct" or "fitting" if and only if *p* is true, considerations independent of whether *p* is true do not constitute normative reasons to believe that *p*. Even if there are practical reasons to believe, they are not the "right kind" of reasons that carry with them the kind of genuine normative force that epistemic reasons do.

While I think that the distinction between content-related and attitude-related reasons is a legitimate and useful one, the "wrong kind" strategy employed here fails, generally speaking, and we should be worried about it as support for epistemic exclusivism. Leary (pgs. 532-3) points this out quite nicely: Not all items of normative evaluation with constitutive correctness conditions, including some action-types, rule out the relevance of reasons that are not content-related. Consider one of her cases: assertion. By many accounts, the content of an assertion that p is correct if and only if p is true. Does that now mean that any consideration that does not count in favor of the truth of p is to be automatically ruled out as a normative reason to assert that p? Surely not! In fact, it is sometimes the case that we have good normative reasons to assert that p, even if we *know* that p is false. If telling a lie could save someone's life, for instance, that fact gives me good normative reason to lie even if saving a life does not count in favor of the truth of the content of the lie.

Similarly, with respect to beliefs, we should not take the attitude-related nature of non-epistemic reasons to automatically rule them out as genuinely normative. Think of beliefs, as many do, as analogous to maps.[144] Beliefs, assertions, and maps, alike, are all representational and have what might be called a "content-to-world" direction of fit. Like a belief, a map is correct if and only if the contents of the map accurately represent reality. Yet, we might have any number of normative reasons to make or draw a particular token map independently of that map's correctness. In fact, we might even have good normative reasons to draw an *incorrect* map. Perhaps we wish to play a joke on a friend, or to mislead some criminals, or to create an ironic work of art. Why not

---

[144] See, e.g., Lewis (1994), Camp (2007), and Rescorla (2009) for arguments that our cognitive representations have map-like structure.

think of belief in the very same way? As with maps, while we often rightly care about accuracy, we may sometimes have good reasons to form beliefs independently of whether such beliefs would be accurate or not. Since we cannot simply and automatically rule out these kinds of state-related reasons, a more generally applicable framework is needed.

It actually seems to me that the traditional "wrong kinds" framework for thinking about reasons is the reverse of what it should be, at least insofar as we are thinking about *guidance-giving reasons* for belief. Recall that content-related reasons are most directly related to the *contents* of the relevant attitudes or states and whether such contents are "correct" or "fitting." For beliefs, content-related reasons are considerations that count in favor of the truth of a target proposition. Given the nature of these reasons, the existence of a content-related reason to believe that *p*, by itself, does not say anything *directly* about the normative constraints surrounding the belief that *p* or whether the belief is to be taken up or not. It merely tells us that there is a consideration that counts in favor of *p*'s truth. Attitude-related reasons, however, when present, *are* directly applicable to the uptake (or maintenance) of the states themselves. When we have an attitude-related reason to believe that *p*, we have a consideration that counts in favor of *believing that p*. At least on the surface of things, attitude-related reasons should be thought of as the directly guidance-giving normative constraints while content-related reasons play a more indirect role.

The implication is that, in order for epistemic reasons (which are content-related) to play some kind of guidance-giving normative role, there would need to be bridging principles between the possession of epistemic reasons and the possession of guidance-giving, attitude-related reasons. Consider a possibility for a bridge principle: If *S* has

good epistemic reason to believe that *p*, then that gives *S* a (more or less strong) attitude-related reason to believe that *p*. Given the truth of such a principle, good epistemic reasons would then provide a guidance-giving normative constraint on *S*'s beliefs. Consider another possibility: If *S* has neither good epistemic reason to believe that *p* nor good epistemic reason to believe that ¬*p*, then that gives *S* a (more or less strong) attitude-related reason to suspend belief as to whether *p*. Whatever the most plausible bridge principles end up being, they are ways in which antecedently aggregated epistemic reasons can generate guidance-giving normative constraints on belief. Presumably, that is theoretically desirable.

The point here, however, is that, even if the epistemic considerations do end up generating attitude-related reasons through such bridge principles, we cannot rule out the possibility that there might be *other* attitude-related reasons that would be in conflict with those epistemically-generated guidance-giving reasons, perhaps even outweighing them in an all-things-considered aggregation about what we ought to believe. We cannot rule this possibility out unless, of course, there is some further argument that could establish it.

This framework in which content-related reasons generate a defeasible class of genuine, guidance-giving reasons does have a more general kind of plausibility to it. Consider, once again, assertion: Suppose that I am a biology teacher and that, in a given moment in class, I am considering whether or not I ought to assert that humans and chimpanzees have a common ancestor. Since the overall evidence counts in favor of the truth of that proposition (along with the fact that it is part of my job to tell my students the truth about biology), a relatively strong guidance-giving reason to make the assertion

that humans and chimpanzees have a common ancestor might be generated. However, is it now the case that this reason is *exclusive* and cannot possibly be overridden by other guidance-giving considerations? Seemingly not. If I am a teacher in a community of fundamentalist zealots who would likely imprison me (or worse) for asserting such things, I might have *more* guidance-giving reason to refrain from making that assertion. Similarly, even if content-related reasons to believe that *p* generate a strong guidance-giving reason to believe that *p*, there is no guarantee that such a guidance-giving reason will not be overridden by other considerations in an all-things-considered aggregation.

Further, consider the fact that, if epistemic reasons are constitutively content-related and they only play a role in guidance-giving through further bridge principles, then, unlike guidance-giving reasons, epistemic reasons can uncontroversially apply independently of whether or not we have a relevant control over our beliefs and independently of whether or not a categorical form of NRP is true. If *S*'s possession of an epistemic reason to believe that *p*, in and of itself, merely tells us that *S* possesses a consideration that counts in favor of the truth of *p* (e.g., in the form of evidence), then *S* can possess that kind of reason in a world in which she does not have the relevant kind of control over her beliefs. There would be no "reason implies can" principle applying to epistemic reasons.

The ability to draw further guidance-giving reasons from those epistemic reasons, however, might depend on the kind of control that *S* has. If *S* does not have a relevant control over her beliefs, then any bridge principle from her epistemic reasons to guidance-giving reasons would be ruled out by proponents of the deliberative constraint on reasons (e.g., people like Shah). This, however, seems to be exactly the right verdict.

We do not need a relevant control over our beliefs in order for us to possess evidence for them (or not), but, given the deliberative constraint, we do need control over beliefs in order for guidance-giving norms to apply directly to them (see Section 4.2 above).

If it is right to frame the distinction between content-related and attitude-related reasons in the way that I have above, there is a possible middle ground position in the exchange between Thomas Kelly and Adam Leite about the character of epistemic reasons.[145] Their exchange concerns whether epistemic reasons are rightly seen as a kind of instrumental reason. Kelly argues that epistemic reasons to believe are not merely instrumental reasons while Leite attempts to make room for the possibility by countering Kelly's arguments. On one hand, the framework above suggests that Kelly is right. An epistemic reason to believe that *p* is *not* merely a kind of instrumental reason. Since instrumental reasons are grounded by psychologically instantiated desires or goals, epistemic reasons are not reasons of this type. Independently of whether *S* has a particular desire about the truth of her own beliefs or not, insofar as *S* possesses evidence (broadly-speaking) for *p*, *S* has an epistemic reason to believe that *p*. *S* possesses a consideration that counts in favor of the truth of *p*, regardless of whether she cares about the truth of *p* or not. In this way, the framework above respects Kelly's requirement about the "intersubjectivity of epistemic reasons." (2003, pg. 621) The same people in the same evidential situations, broadly-speaking, have the same epistemic reasons for belief.

On the other hand, Leite is right in that we need to keep the question of what we *ought* to believe separate from the question of what our evidence is or of what our

---

[145] See Kelly (2003; 2007) and Leite (2007).

epistemic reasons are, at least insofar as we are talking about the guidance-giving ought. If the distinction given above is legitimate, good epistemic reasons and guidance-giving oughts should not be conflated and it is an open question of what the bridge principles between them will look like (if there are any at all). We cannot just assume that the guidance-giving rationality of belief will neatly fall in line with epistemic rationality.[146]

Some authors have thought that, when our epistemic reasons conflict with our practical reasons for belief, there is just no further question about what we ought to believe. We cannot simply weigh each of the considerations up against one another and come to an overall normative verdict.[147] The above would suggest that there is *something* right about this. For instance, given the distinction between attitude-related and content-related reasons, any attempt to directly aggregate epistemic reasons and practical reasons for belief, any attempt at all to weigh them up against one another, is a fool's errand. Attitude and content-related reasons are different kinds of reasons; they are relations holding between fundamentally different kinds of things, the former favoring the uptake (or maintenance) of an attitude or state and the latter favoring the truth of a proposition.[148]

Yet, this is all consistent with the idea that there are facts of the matter about what we guidance-giving ought to believe, even in cases where one's practical reasons conflict with what would be epistemically supported. What we ought to believe in the guidance-giving sense ultimately depends on the correct normative bridge principles connecting

---

[146] Here, I use the term "rationality of belief" to refer to the state of affairs of a belief being held for the right reasons, not merely the state of a belief when it meets certain consistency requirements.

[147] See, e.g., Feldman (1988, 2000) and Engel (2013).

[148] I think that this is the best explanation for the frequently-cited fact that practical reasons exhibit what Selim Berker (unpublished manuscript) calls "permissive balancing" while epistemic reasons exhibit "prohibitive balancing" (also see, e.g., Harman 1995, Feldman 2000, and Dancy 2004). I discuss Berker's worries in more detail later in this chapter.

epistemic rationality and our guidance-giving, attitude-related reasons and on how those

attitude-related reasons weigh up against other kinds of attitude-related reasons. The

plausibility of a particular set of bridge principles will depend on many things, including

complex meta-normative issues that arise in formulating a broader, overall theory of

guidance-giving reasons. The plausibility of a set of principles will ultimately depend on,

for instance, whether a so-called "Humean approach" to guidance-giving reasons is

adequate or not (i.e., whether or not we can reduce facts about guidance-giving reasons to

facts about our pro-attitudes).[149] They will depend on how reasons are related to value,

whether a "value-first" theory of reasons is adequate or whether something more like a

"buck-passing" account of value is true, among many other things.

We should survey at least some of the possibilities here and how an epistemic

exclusivist might wish to proceed. One might want to propose an exclusivist bridge

principle that falls in line with SCT above (see pg. 152); call it the "Simple Connection

Thesis for Reasons" (or "SCTR"): For all S and for all p, if it would be epistemically

rational for S to believe that p, then S has most guidance-giving reason to believe that p

(and similarly for disbelieving that p and suspending belief as to whether p).[150] SCTR,

again, would rule out the normative relevance of non-epistemic considerations. But,

what kind of support could an exclusivist give it?

For the sake of argument, let us adopt a meta-normative narrative about reasons

and see what can be done. Suppose, first, that considerations about *value* ground facts

---

[149] For an example of a Humean approach, see Schroeder (2007).

[150] SCTR is at least extensionally equivalent to SCT insofar as a traditional kind of connection between epistemic justification and epistemic reasons holds (i.e., that being epistemically justified in some sense amounts to being epistemically rational).

about guidance-giving reasons.[151] The "value-first" idea is something like the following: The fact that S's φ-ing would promote (or would justifiably promote) some valuable state of affairs is a guidance-giving reason for S to φ.[152] Similarly, the fact that S's φ-ing would promote (or would justifiably promote) some *dis*valuable state of affairs is a guidance-giving reason for S *not* to φ. In determining whether it is guidance-giving rational to φ or not, one must aggregate these value-based reasons for and against φ-ing (along with the value-based reasons for and against the alternatives to φ-ing) and relevantly compare. In such a framework, the totality of the guidance-giving reasons in a given situation will depend on a particular substantive axiology, on the states of affairs that are ultimately valuable and disvaluable. Is there some plausible story about the *value* of epistemic rationality and the aggregation of the related guidance-giving reasons that could push us toward epistemic exclusivism?

There are a number of reasons to doubt that there is such a story. One reason to doubt is that the mere *epistemic* value of epistemic rationality (or justification) and the states of affairs in which it obtains might not be enough to give it a relevant role in automatically generating guidance-giving reasons. Recall that, if the critical domain approach to epistemic evaluation developed in the first three chapters has any merit, then epistemic value might be a kind of *domain-internal* value. Hence, for some particular epistemically valuable state of affairs, it would still be a substantive question as to whether any *genuine* value obtains (see pgs. 11-5 in Chapter 1). The fact that some state

---

[151] For examples of such an approach, see Steglich-Petersen (2011) and Maguire (2016). Scanlon (1998, pg. 84) also discusses an approach like this, though he ultimately rejects it.

[152] "Promote," here, is intended in a broad sense. Maguire, in particular, takes φ-ing to "promote" some state of affairs if it instantiates the state of affairs, causes (or partially causes) it, constitutes (or partially constitutes) it, prevents the preventing of it, or non-superfluously probabilifies it. (pg. 238)

of affairs is domain-internally valuable does not entail that it is genuinely valuable and, presumably, it is genuine value that ultimately grounds guidance-giving reasons in a value-first framework. As to the more substantive axiological questions, I have also argued that there are items of *fundamental* epistemic value (i.e., instances of knowledge) that do not have any real, genuine value at all (see pgs. 57-8 in Chapter 2). This should also raise questions about the idea that each and every instance of epistemic rationality is genuinely valuable.

But let us leave these issues aside and assume, for the moment, that instances of epistemic rationality are always genuinely valuable and that they do always generate guidance-giving reasons in the way a value-first evidentialist would have it. The issue now is that these generated reasons seem to be no more than *pro tanto* reasons, defeasible by other axiological considerations. In order to generate SCTR, we would need to guarantee that the guidance-giving reasons produced by the value of epistemic rationality are never defeated by reasons produced by other kinds of value. At least on the surface of things, it seems highly implausible that this could be established. Consider Christie, once again: Even if it was the case that there is some genuine value in her having an epistemically rational doxastic attitude toward whether God exists and that this gives her a guidance-giving reason to take on an epistemically rational doxastic attitude, why would it be impossible for such a reason to be overridden by other guidance-giving reasons generated by other, non-epistemic axiological considerations?

We should be skeptical of the idea that it is *often* more valuable to hold an epistemically irrational doxastic attitude than it is to hold an epistemically rational one, but SCTR requires more than this. If we take the most natural way of developing a

value-first framework where the strength of one's guidance-giving reasons for or against

φ-ing is proportional to the overall amount of value promoted by φ-ing, then SCTR

would require that taking on an epistemically rational doxastic attitude is *always* more

valuable overall than taking on an epistemically irrational one. In order to establish this,

we would need to be point out some special feature of the value of epistemic rationality

that could explain why instances of it are always more valuable than the alternatives.

Is there such a feature? Hilary Kornblith (1993) attempts a value-first explanation

of the universal normative force of epistemic evaluation. According to him, epistemic

evaluation is unique because epistemically sanctioned beliefs are reliably accurate and

having accurate beliefs is a necessary part of promoting other valuable states of affairs,

whatever those states of affairs might be. Kornblith claims that, even if value turns out to

be a relative notion, when we are deciding how to act or believe, we must each still

perform cost-benefit calculations in order to determine how to properly promote that

value. In order for those cost-benefit calculations to assist us in any real value

promotion, the outputs of those calculations must *accurately* represent the relations

between our actions/beliefs and the valuable states of affairs that we aim to promote. He

summarizes his argument:

> This suggests that epistemic evaluation takes on a special role… Precisely
> because our cognitive systems are required to perform evaluations relative to our
> many concerns, and to perform these evaluations accurately, the standards by
> which we evaluate these cognitive systems themselves must remain insulated
> from most of what we intrinsically value, whatever we may value. This provides
> a reason to care about the truth whatever we may otherwise care about. (pg. 372)

The idea seems to be that, if there are any valuable states of affairs requiring promotion at

all, since we need to accurately gauge those value promotion relations, the norms

governing our beliefs must be universally aimed toward truth. Hence, epistemic

evaluation has a special kind of normative force, no matter what turns out to be valuable.

Could we use such a feature of epistemic evaluation in support of SCTR? There

are many reasons to think that Kornblith is right about the importance of accuracy in our

cognitive systems. After all, our perceptual and cognitive systems have evolved together

in order to generate a relatively reliable access to our environments. If we had perceptual

systems that represented the world inaccurately, leading us to have false beliefs about our

environments, we would take that as a devastating kind of impairment, as severely

detracting from our ability to lead a healthy life. And this is true for more than the mere

perceptual inputs. If, in general, the way we reasoned about the world led us to massive

error, we would not be able to achieve much of anything at all. Take a very simple

example: If I needed to take an antidote after consuming some poison, I would be pretty

helpless if I falsely believed that the antidote would worsen my condition or if I falsely

believed that I did not consume any poison to begin with.

Yet, even if, on the whole, we want our cognitive systems to be accurate for those

reasons, there still seems to be an open question about whether a *universal* guidance-

giving norm follows, whether it is everywhere and always, no matter what the belief is,

normatively binding upon us to believe in an epistemically rational way. There is a

difference between the betterness or worseness of an individual's overall cognitive

system and the betterness or worseness of a token belief at a particular time. Perhaps

more relevantly to Kornblith's argument, there is a difference between the value of an

accurate belief that *p* and the value of an accurate belief about whether that original belief

promotes value or not. Applying Kornblith's point, independently of whether an accurate

belief that *p* is valuable or not, it is still valuable for us to have an accurate gauge of that belief's value-promotion relations and, hence, we still have good guidance-giving reason to care about accuracy. But, even so, this only tells us that it is valuable to have accurate beliefs about the value-promotion relations relevant to the original belief that *p*. It tells us nothing, in particular, about the value of the original belief that *p* or the norms governing it. Everything that Kornblith says is consistent with some particular accurate belief that *p* being all-things-considered *disvaluable* even if it is all-things-considered valuable to have accurate beliefs about the relevant value-promotion relations.

Perhaps an application will bring this point out more clearly. There have been a number of psychological studies correlating a kind of self-enhancement bias (i.e., having overly positive beliefs about one's own traits or abilities) with rates of non-depression and higher self-esteem.[153] If those studies are legitimate and we think of those beliefs as a species of epistemic irrationality, then, at least in some cases, it seems as though these epistemically irrational beliefs can manifest a genuine kind of value. Given a value-first approach to guidance-giving reasons, why should we think that the reasons generated by such axiological considerations are *always* overridden by the sorts of reasons generated by the value of believing in an epistemically rational way? To Kornblith's point, even if it is all-things-considered valuable to have accurate beliefs about whether or not those overly inflated beliefs promote value or not, this tells us nothing in particular about the actual value of those overly inflated beliefs.

---

[153] For some general reviews of these studies, see Taylor and Brown (1988) and Alloy and Abramson (1988). Also see Hazlett (2013, Ch. 2) for a general discussion of the axiological implications of self-enhancement bias.

Could we guarantee that *none* of those overly inflated beliefs are all-things-considered most valuable as compared to the relevant doxastic alternatives? Take Adam, a pretty generally low-skilled individual, who has some overly inflated beliefs about his abilities with respect to a particular hobby. Adam does not look forward to doing much in his life but he really loves his hobby and his beliefs about his hobby-related skills are able to help him cope and they boost his self-esteem. Even if it would be valuable for Adam to have accurate beliefs about how his hobby-related confidence improves his own psychology and self-esteem, would it really be plausible to think that his epistemically irrational self-confidence is less valuable than the more epistemically justified doxastic alternatives and that he thereby ought to give those beliefs up? What would be the basis for such a claim?

One might think with Clifford (1947) that, when we allow even a single epistemically unjustified belief into a cognitive system, the epistemic irrationality involved "spreads" in a way that is detrimental to functionality of the entire system. Hence, focusing on the kind of practical value generated by boosts in self-esteem and the like is short-sighted way of evaluating something that is ultimately harmful. Clifford claims: "Every time we let ourselves believe for unworthy [epistemic] reasons, we weaken our powers of self-control, of doubting, of judicially and fairly weighing the evidence." (pg. 76) Purportedly, when we do receive practical benefits from believing in an epistemically unjustified way, those benefits are still always outweighed by the detrimental effects to our overall cognitive systems because those beliefs ultimately

undermine our ability to form and maintain other beliefs where accuracy might be more important.[154]

Of course, however, the actual psychological effects of self-enhancement bias (or of any other kind of epistemic irrationality) are an empirical matter and, if there are such universal detrimental effects, we would expect them, at least to a significant degree, to be reflected in the empirical research. As a matter of fact, however, those effects are not substantiated. As Allan Hazlett (2013, pgs. 41-9 and 80-2) makes clear, the empirical research indicates that we quite often compartmentalize our epistemic irrationality, insulating such beliefs and the doxastic practices that sustain them from the rest of our "web of belief," protecting ourselves from massive delusion. Moreover, McKay and Dennett (2009) compellingly argue that a number of our epistemically irrational beliefs (i.e., our "misbeliefs" as they call them), including those manifested in self-enhancement bias, are actually adaptive, evolutionarily selected for the propagation of the species. Given all of the empirical evidence, it is just not clear at all that epistemically irrational beliefs are *always* all-things-considered harmful, individually or socially, in a way that could guarantee that the relevant guidance-giving reasons would be overridden. For all of these reasons, among others, we should be skeptical that a value-first framework for guidance-giving reasons can plausibly generate anything like SCTR.

For the purposes here, it is just as important to note that even if a value-first framework *could* do this, it would do so in a way that is perfectly consistent with the critical domain theory of epistemic evaluation outlined above. Insofar as there are facts *extrinsic* to the epistemic domain itself that are supposed explain the appropriateness of

---

[154] Paul Horwich (2006) argues in a similar way: The truth of a belief is always crucial because "there is no proposition that might not someday serve as a premise." (pg. 350)

epistemic exclusivity, the critical domain theorist is free to appeal to such facts. We have been considering whether there is a reasonable story to tell about the *genuine value* of believing in epistemically sanctioned ways that could explain the universal normative force of epistemic evaluation. The critical domain theorist who takes epistemic value to be *domain-internal*, grounded by the constitutive goals of inquiry, is still free to further her story in this way. Though we have seen a number of reasons to be skeptical about that narrative, the general meta-normative approach is not itself problematic.

Something similar can be said of more "deontological" meta-normative approaches. Suppose, for example, that we could successfully argue that epistemically irrational belief manifests a kind of self-disrespect or self-contempt and it is just a normative rule that we ought not disrespect ourselves in such ways (see, e.g., Wood 2008, pgs. 18-20). Dialectical effectiveness aside, given that there *is* good reason to be a deontologist of this kind and insofar as appealing to such facts does provide a strong basis for some kind of epistemic exclusivism, the critical domain theorist is free to agree. Normative facts like these are just not inconsistent with a critical domain approach to epistemic evaluation.[155] Similarly, if it turns out that we have most reason to believe that SCTR is just some kind of primitive fact, that it is just an inexplicable truth that one always has most guidance-giving reason to be epistemically rational. The truth of epistemic exclusivism really only constitutes a barrier for the critical domain theorist

---

[155] This is also the way I am thinking of Chrisman's (2008) approach to doxastic oughts. According to Chrisman, true claims about what we ought to believe are "ought-to-be" norms issuing from certain human ideals. In the case of doxastic oughts, the relevant ideals are truth related and are purportedly categorical, grounded in our natures as "tracking and transmitting beings." (pgs. 366-7) Perhaps it is correct that believing the truth is a human ideal in a way that could generate categorical "ought-to-be" norms about our beliefs. However, if it is true, once again, I see no barrier to the critical domain theorist furthering her story in this way. Dialectical effectiveness aside, whatever facts are supposed to make it the case that believing the truth is a categorical human ideal, those facts are not necessarily inconsistent with the view presented in these chapters.

insofar as we demand that there must be facts *intrinsic to the epistemic domain itself* that could explain its truth, that it is something constitutive of epistemic evaluation itself that can establish its exclusivity in the norms governing belief. But, given the distinctions that have been made above, I do not know why such a demand would be an appropriate one.

I will discuss one, final worry here. The basic framework that has been set up in this section incorporates epistemic reasons as content-related reasons, specifically as considerations that count in favor of the truth of some proposition. I argued that, in order for such reasons to play any role in guidance-giving normativity, certain bridge principles must be established between the epistemic rationality of a doxastic state and the guidance-giving reasons relevant to that state. If it is established that some particular doxastic state is an epistemically rational one through the appropriate aggregation and comparison of epistemic reasons (making the relevant doxastic alternatives irrational), then some kind of bridge principle is needed in order to generate guidance-giving reasons counting in favor of that state (or to generate reasons counting against the alternative states). Depending on what those bridge principles look like, those epistemically-generated guidance-giving reasons might be more or less strong, more or less likely to override competing guidance-giving considerations when aggregated and compared.

This framework is relevantly similar to what Berker (unpublished manuscript) calls "double weighing" versions of "austere pragmatism" with respect to reasons for belief. Double weighing views, according to Berker, require that the epistemic considerations relevant to a particular proposition be aggregated and compared before the result can, in some way or other (via what I have called "bridge principles"), feed into the

aggregation and comparison of the ultimately normative reasons for belief. There are two significant differences between the framework set out here and this kind of "austere pragmatism:" (1) While austere pragmatism is, appropriately, an explicit form of pragmatism, the framework that I have set up here is ultimately non-committal in the evidentialism/pragmatism debate. While the framework certainly leaves pragmatism as an open possibility, it is also an open question whether something like SCTR is a legitimate bridge principle and I have only argued that, *if* evidentialism turns out to be true, then it is consistent with a critical domain approach to epistemic evaluation. (2) Berker's "austere pragmatism" rejects the existence of epistemic reasons. I have not denied the existence of epistemic reasons; in fact, I have affirmed their existence. I have only argued that such reasons are not constitutively guidance-giving.

Neither of these differences, however, is crucial to a particular objection that Berker brings out against austere pragmatism. His objection would apply equally against the framework suggested here. Berker argues that double weighing procedures are bound to be problematic and cannot be a proper model for thinking about reasons for belief. If this would indeed be the case, then the arguments above would have little to recommend themselves. Berker formulates a specific example that is supposed to be problematic for double weighing views. We are to imagine that, through time, though I steadfastly disbelieve that *p*, my situation with respect to *p* changes in a way that I transition from (i) to (ii) to (iii):

(i)   Something terrible will happen if I believe that *p*, and I have strong *pro tanto* evidence against *p* (and no other evidence that bears on the matter).

(ii)   Something terrible will happen if I believe that *p*, and I have no *pro tanto* evidence for or against *p*.

(iii) Something terrible will happen if I believe that *p*, and I have strong *pro tanto* evidence for *p* (and no other evidence that bears on the matter).

Berker claims that the double weighing austere pragmatist is committed to the verdict

that, while my evidential situation towards p steadily improves over time, my attitude of

disbelieving that p goes from being permitted, to being forbidden, to being permitted

again, and that this is an odd verdict to be committed to, especially in the final transition.

As he claims, "How can my disbelief in p go from being forbidden to being permitted…

when the only relevant change is that I have gained some evidence for p?" (pg. 27)

While Berker does not explain in detail his claim that any double weighing

austere pragmatist would be committed to such a verdict (he mostly relies on his

treatment of a proposal from Andrew Reisner), it is worth attempting a rational

reconstruction here if only to contrast it with what I will eventually propose. Presumably,

Berker thinks that disbelieving that p in (i) would be at least permitted by the double

weighing austere pragmatist because, insofar as the epistemic rationality of a state

generates some *pro tanto* normative reason in favor of it (and the epistemic irrationality

of a state generates some *pro tanto* normative reason against it), disbelieving that p is

what I have all-things-considered most normative reason to do in (i). In (i), while I have

some reason to disbelieve that p (because that would be the epistemically rational option),

I have only reasons against suspending belief about whether p (because that would be

epistemically irrational) and believing that p (because that would be both epistemically

irrational and practically disastrous). In (ii), disbelieving that p becomes forbidden,

presumably because I now have most all-things-considered normative reason to suspend

belief about whether p. After all, suspending is the epistemically rational option (as

opposed to my disbelief, which is epistemically irrational) and I still only have reasons

against believing that p. In (iii), even though believing that p is now the epistemically

rational option (giving us some *pro tanto* normative reason to believe that *p*), we are supposed to imagine that this reason is far outweighed by the disastrous practical results of believing that *p*, still giving me very strong all-things-considered reason *not* to believe that *p*. My disbelief is supposed to be permitted because, while both disbelief and suspension are epistemically irrational, the normative reasons against disbelieving and suspending are not as strong as the normative reasons against believing.

I will show here that a proponent of the framework outlined above is not necessarily committed to a permissible-forbidden-permissible verdict in the transition from (i) to (iii) and that, when she *is* so committed, there is perfectly good explanation available as to why it is appropriate. Berker is too quick to assume that a double weighing austere pragmatist must always treat the transition from (i) to (iii) in this way. This is partly because, like the ways in which the strength of normal practical reasons are sensitive to the *importance* of the relevant practical results, the bridge principles connecting epistemic rationality to guidance-giving reasons might also be sensitive to the *importance* of epistemic rationality given the circumstances. For a double weighing view, the most plausible way of connecting epistemic rationality to guidance-giving reasons would be to make the strength of the relevant epistemically-generated guidance-giving reasons proportional to the level of importance of epistemic rationality given the specific circumstances.[156] When it is more important to be epistemically rational in a given situation, a stronger guidance-giving reason to be epistemically rational is generated (and similarly for reasons *against* adopting epistemically irrational states).

---

[156] This is not necessarily to commit to the view that such reasons are *grounded* by those considerations about importance or value or anything else of the sort. It is merely to commit to an extensional coincidence thesis, to the strength of our reasons mirroring the level of importance of the relevant considerations.

When it is less important to be epistemically rational, a weaker guidance-giving reason to be epistemically rational is generated.

Consider, now, the transition from (i) to (iii) under the assumption that the epistemic rationality of my doxastic state toward p is *not* a particularly important state of affairs given my situation, that there is only a very small value or positive significance in being epistemically rational with respect to p and only a very small disvalue or negative significance in being epistemically irrational. Should the verdict regarding my disbelief be permissible, then forbidden, then permissible again? Consider (i): Given that I have even some small guidance-giving reason to be epistemically rational (and thereby to disbelieve that p) and that I have only reasons *not* to suspend belief (because it would be epistemically irrational) and *not* to believe that p (because it would be epistemically irrational and practically disastrous), it seems as though I have *most reason* to disbelieve that p. Hence, I *ought* to disbelieve that p (see pgs. 157-8) and, insofar as it is permissible to do what I ought to do, permissibility is an accurate verdict.

Yet, while it is appropriate to say that my disbelief in (i) is rationally permissible, I am also inclined to say that suspending belief would have been permissible as well. While I do not have *most reason* to suspend belief in (i), we have stipulated that the epistemic rationality of my doxastic state with respect to p is not an important factor, meaning that the differences in the guidance-giving reasons generated by considerations about epistemic rationality are very small. In that case, the ultimate difference in the force of the reasons favoring disbelief and the reasons disfavoring suspension in (i) are very small.

In support of this, consider a plausible way of thinking about permissibility: The rational permissibility of some option φ for S at t is a comparative function of the difference between the strength of the reasons relevant to S's φ-ing at t and the strength of the reasons relevant to what S ought to do at t. More specifically, if the difference between the strength of the reasons relevant to φ-ing and the reasons relevant to what one ought to do is small to a relevant degree, then φ-ing is permissible even if φ-ing is not what one ought to do. Suppose, for example, that you are forced to choose between 3 different options for dinner and that, while you have very strong reasons *not* to choose option 1 (suppose that it contains large quantities of a deadly poison), there are only negligible differences between the strength of the reasons with respect to choosing options 2 and 3 (suppose that there are very tiny differences related to the healthiness of the foods). I think the appropriate verdict here is that, while choosing option 1 would be rationally forbidden (in Berker's terms, you have "decisive reason" not to choose it), choosing either option 2 or option 3 would be rationally permissible even if you have slightly more reason to choose one over the other. The differences in the strength of the reasons relevant to options 2 and 3 are small enough to make either choice rationally benign.[157] The point here is that I am in an analogous position in (i). While it would be rationally forbidden to believe that p, it would be permissible to either suspend belief as to whether p or to disbelieve that p given that the differences in the strength of the reasons between the two options are negligible.

Consider, now, the transition from (i) to (ii). Berker claims that, from the double weighing austere pragmatist's perspective, disbelieving that p goes from being

---

[157] For some agreement here, see Slote (1989). In a similar way, Slote claims that some option can be less than ideally rational but yet rationally permissible.

permissible to being forbidden. If the above is correct, however, this is not the right verdict. While believing that *p* is still going to be rationally forbidden in (ii) because of the disastrous practical consequences (combined with the epistemic irrationality) and while I now have *most reason* to suspend belief as to whether *p* because the epistemically-generated guidance-giving reasons have relevantly shifted, there is still only a negligible difference between the strength of the reasons relevant to suspending belief and the reasons relevant to disbelieving. This would indicate that it would still be permissible to disbelieve that *p*. But, I do not think that this is a strange verdict. We have stipulated from the beginning that the epistemic rationality of the doxastic state that I take towards *p* is unimportant. It would make sense, then, that the permissibility of disbelieving would not change when only the epistemic considerations change.

We can say something very similar about the transition from (ii) to (iii). In (iii), while it would be epistemically irrational to either suspend belief or to disbelieve that *p*, the generated guidance-giving reasons are still far outweighed by the all-things-considered guidance-giving reasons *not* to believe that *p*. Even though the all-things-considered reasons not to believe that *p* are slightly diminished because I now have a small *pro tanto* reason in favor of believing that *p* in virtue of the epistemic rationality involved, believing that *p* is still rationally forbidden in (iii) because of the disastrous practical consequences. Given the relative unimportance of epistemic rationality in the circumstances, the correct verdict about Berker's transition is permissible-permissible-permissible.

Now, let us consider the transition from (i) to (iii) when the epistemic rationality of my doxastic state toward *p* is *crucially important*. In fact, suppose that the importance

of me being epistemic rational is somewhat comparable to the importance of the supposedly disastrous results of believing that p. Generally speaking, this means that the guidance-giving reasons generated by the epistemic rationality or irrationality of the various options will be much stronger. In (i), disbelieving that p is still what I rationally ought to do and, hence, it would be permissible. However, the difference now is that, in (i), suspending belief would no longer also be permissible. Since we have stipulated that it is much more important to be epistemically rational, the differences in the strength of the reasons relevant to suspending and to disbelieving is no longer negligible. In (i), I am rationally forbidden to suspend belief because it would be epistemically irrational and that irrationality generates a very strong and significant guidance-giving reason not to suspend. When we transition to (ii), the result is that my disbelief is rationally forbidden because the roles of suspending and disbelieving are reversed. So far, this is what Berker expects.

The problem is that, when we transition to (iii), disbelieving that p does not go back to being permissible as Berker claims. In (iii), suspending and disbelieving are rationally *forbidden* because believing that p is significantly rationally superior to both. While the disastrous practical results of believing that p do generate a strong *pro tanto* reason not to believe that p, that *pro tanto* reason is countered by the *pro tanto* reason generated by the fact that believing that p is the epistemically rational doxastic state. Given the stipulated importance of epistemic rationality, the reason against believing that p is effectively cancelled. Hence, in comparison, the net balance of reasons relevant to believing makes that option rationally superior in a significant way to both suspending and disbelieving. Under the circumstances, the correct verdict is then permissible-

forbidden-forbidden. This is not a strange verdict given the importance of epistemic

rationality and given that disbelieving goes from being epistemically rational to being

epistemically irrational.

The only time in which Berker's alleged verdict is correct is when we assume that

it is *moderately* important to a very specific degree for me to take on an epistemically

rational state toward *p*. His verdict becomes appropriate because, in the transition to (iii),

while the *pro tanto* practical reason against believing that *p* is still countered by an

epistemically-generated reason, it is not effectively cancelled out and there is a

remaining, all-things-considered reason *not* to believe that *p*. Given the appropriate

weights, this all-things-considered reason not to believe that *p* can be made comparable to

the reasons not to suspend and not to disbelieve, respectively. Hence, in (iii), any of the

three options would be permissible. Since (i) and (ii) should be treated no differently

than above, Berker's alleged verdict of permissible-forbidden-permissible is the right

one.

Even so, I think that we have the right tools to explain this verdict and that it is

not a strange one in any harmful way. Its strangeness is merely a benign product of the

gerrymandered circumstances. Consider it this way: Given that epistemic rationality with

respect to *p* is moderately important, the verdict in the transition from (i) to (ii) is not at

all surprising. After all, the epistemic rationality of the state is moderately important and

my disbelieving goes from being epistemically rational to being epistemically irrational

while all other things remain equal. It should not be shocking that my disbelief goes

from being permissible to forbidden. In (ii), there is now a *clearly better option* available

to me: suspension. This was not the case in (i).

In the transition from (ii) to (iii), my disbelief goes from being forbidden to permissible because I go from having a clearly better option to having *no good option at all*. Given the relevant weights of the reasons, my options are all *equally bad* in (iii) and, when our options are equally bad, we may permissibly choose any of them (precisely because no option is clearly better than any other). While this transition might be a strange one, it is strange in an explicable way. Hence, Berker's case does not pose any threat to what has been presented in this section.

To summarize the details of this chapter, I have considered some ways in which considerations about the ethics of belief may present problems for the critical domain theory of epistemic evaluation presented above. At least on the surface of things, it may seem as though a kind of *evidentialist* position in the ethics of belief is inconsistent with the view of epistemic justification presented in Chapter 3 and with the goals of the project more generally. I have argued, however, that certain ways of arguing for evidentialism (via NRP) are bound to fail and are thus not capable of challenging the critical domain picture. I have also argued that alternative routes to evidentialism are either dialectically ineffective (e.g., those utilizing the distinction between the so-called "right" and "wrong" kinds of reasons) or are forms of evidentialism that would be completely consistent with seeing epistemic evaluation as a critical domain framed by the constitutive goals of inquiry. On the whole, this project gives us a way of looking at the big picture concerning epistemic evaluation, the tools for explaining what it is, why it is important for us, and how it plays into our lives as creatures bound by norms.

List of References

Achinstein, Peter (1983). *The Nature of Explanation*. Oxford, Oxford University Press.

Adler, Jonathan (2002). *Belief's Own Ethics*. Cambridge, MIT Press.

Adler, Jonathan and Hicks, Michael (2013). "Non-Evidential Reasons to Believe," in *The Aim of Belief*. Ed. Timothy Chan. Oxford, Oxford University Press.

Ahlstrom-Vij, Kristoffer (2012). In Defense of Veritistic Value Monism. *Pacific Philosophical Quarterly*. 94(1): 19-40.

Ahlstrom-Vij, Kristoffer and Grimm, Stephen (2013). Getting it Right. *Philosophical Studies*. 166(2): 329-47.

Alloy, L. B. and Abramson, L. Y. (1988). "Depressive Realism: Four Theoretical Perspectives," in *Cognitive Processes in Depression*. Ed. L. B. Alloy. New York, Guilford Press.

Alston, William (1983). "Christian Experience and Christian Belief," in *Faith and Rationality*. Eds. Alvin Plantinga and Nicholas Wolterstorff. South Bend, University of Notre Dame Press.

Alston, William (1988a). An Internalist Externalism. *Synthese*. 74(3): 265-83.

Alston, William (1988b). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*. 2: 257-99.

Alston, William (2005). *Beyond "Justification:" Dimensions of Epistemic Evaluation*. Ithaca, Cornell University Press.

Anscombe, G. E. M. (1957). *Intention*. Oxford, Blackwell Publishing.

Aristotle (1984). *Complete Works*. Ed. J.Barnes. Princeton, Princeton University Press.

Armstrong, David (1973). *Belief, Truth, and Knowledge*. Cambridge, Cambridge University Press.

Audi, Robert (2001). "Epistemic Virtue and Justified Belief," in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford, Oxford University Press.

Baehr, Jason (2009). "Is There a Value Problem?" in *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Bennett, Jonathan (1990). Why Belief is Involuntary? *Analysis*. 50(2): 87-107.

Berker, Selim (2013a). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review*. 122(3): 337-93.

Berker, Selim (2013b). The Rejection of Epistemic Consequentialism. *Philosophical Issues*. 23(1): 363-87.

Berker, Selim (2015). Reply to Goldman: Cutting Up the One to Save the Five in Epistemology. *Episteme*. 12(2): 145-53.

Berker, Selim (unpublished manuscript). A Combinatorial Argument against Practical Reasons for Belief. Draft as of July 27, 2017.

Bird, Alexander (2007). Justified Judging. *Philosophy and Phenomenological Research*. 74(1): 91-110.

Blackburn, Simon (2001). "Reason, Virtue, and Knowledge," in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford, Oxford University Press.

BonJour, Laurence (1985). *The Structure of Empirical Knowledge*. Cambridge, Harvard University Press.

Brady, Michael (2009). "Curiosity and the Value of Truth," in *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Brogaard, Berit (2006). Can Virtue Reliabilism Explain the Value of Knowledge? *Canadian Journal of Philosophy*. 36(3): 335-54.

Buckareff, Andrei (2006). Doxastic Decisions and Controlling Belief. *Acta Analytica*. 21(1): 102-14.

Camp, Elisabeth (2007). Thinking with Maps. *Philosophical Perspectives*. 21(1): 145-82.

Carruthers, Peter (1992). *Human Knowledge and Human Nature*. Oxford, Oxford University Press.

Chisholm, Roderick (1989). *Theory of Knowledge*. Third Edition. Englewood Cliffs, Prentice-Hall Publishing.

Chrisman, Matthew (2008). Ought to Believe. *The Journal of Philosophy*. 105(7): 346-70.

Chrisman, Matthew (2012). The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions. *Journal of Philosophy*. 109(10): 588-612.

Christensen, David (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford, Clarendon Press.

Clifford, William (1947). "The Ethics of Belief," in *The Ethics of Belief and Other Essays*. London, Watts and Co.

Conee, Earl (1982). Against Moral Dilemmas. *The Philosophical Review*. 91: 87-97.

Conee, Earl and Feldman, Richard (1985). Evidentialism. *Philosophical Studies*. 48(1): 15-34.

Conee, Earl and Feldman, Richard (1998). The Generality Problem for Reliabilism. *Philosophical Studies*. 89(1): 1-29.

Copp, David (2007). *Morality in a Natural World: Selected Essays in Metaethics*. Cambridge, Cambridge University Press.

Dancy, Jonathan (2004). "Enticing Reasons," in *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Eds. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith. Oxford, Oxford University Press.

Danto, Arthur (1965). Basic Actions. *American Philosophical Quarterly*. 2(2): 141-8.

Darwall, Stephen (1997). "Reasons, Motives, and the Demands of Morality: An Introduction," in *Moral Discourse and Practice*. Eds. Stephen Darwall, Alan Gibbard, and Peter Railton. New York, Oxford University Press.

David, Marian (2001). "Truth as the Epistemic Goal," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Ed. Matthias Steup. Oxford, Oxford University Press.

David, Marian (2014). "Truth as the Primary Epistemic Goal: A Working Hypothesis," in *Contemporary Debates in Epistemology*. Second Edition. Eds. Matthias Steup, John Turri, and Ernest Sosa. West Sussex, Blackwell Publishing.

Davidson, Donald (1987). Knowing One's Own Mind. *Proceedings and Addresses of the American Philosophical Association*. 60: 441-58.

DePaul, Michael (2001). "Value Monism in Epistemology," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Ed. Matthias Steup. Oxford, Oxford University Press.

DePaul, Michael (2009). "Ugly Analyses and Value," in *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Deutsch, Harry (1998). Identity and General Similarity. *Philosophical Perspectives*. 12: 177-99.

Dogramaci, Sinan (2012). Reverse Engineering Epistemic Evaluations. *Philosophy and Phenomenological Research*. 84(3): 513-30.

Douven, Igor (2008). The Lottery Paradox and our Epistemic Goal. *Pacific Philosophical Quarterly*. 89(2): 204-225.

Dretske, Fred (1981). *Knowledge and the Flow of Information*. Cambridge, MIT Press.

Elgin, Catherine (1996). *Considered Judgment*. Princeton, Princeton University Press.

Elgin, Catherine (2004). True Enough. *Philosophical Issues*. 14: 113-31.

Elgin, Catherine (2009). "Is Understanding Factive?" in *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Engel, Pascal (2013). "In Defence of Normativism about the Aim of Belief," in *The Aim of Belief*. Ed. Timothy Chan. Oxford, Oxford University Press.

Ewing, A. C. (1947). *The Definition of Good*. New York, MacMillan.

Feldman, Richard (1988). Epistemic Obligations. *Philosophical Perspectives*. 2(1): 235-56.

Feldman, Richard (2000). The Ethics of Belief. *Philosophy and Phenomenological Research*. 60(3): 667-95.

Feldman, Richard (2001). "Voluntary Belief and Epistemic Evaluation," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Ed. Matthias Steup. Oxford, Oxford University Press.

Feldman, Richard (2002). "Epistemological Duties," in *The Oxford Handbook of Epistemology*. Ed. Paul Moser. New York, Oxford University Press.

Fine, Kit (1994). Essence and Modality. *Philosophical Perspectives*. 8: 1-16.

Foote, Philippa and Montefiore, Alan (1961). Goodness and Choice. *Proceedings of the Aristotelian Society, Supplementary Volumes*. 35: 45-80.

Foley, Richard (1987). *The Theory of Epistemic Rationality*. Cambridge, Harvard University Press.

Foley, Richard (1993). *Working without a Net: A Study of Egocentric Epistemology*. Oxford, Oxford University Press.

Frankish, Keith (2007). Deciding to Believe Again. *Mind*. 116(463): 523-47.

Fumerton, Richard (2001). "Epistemic Justification and Normativity," in *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford, Oxford University Press.

Garber, Daniel (2007). "*Religio Philosophi*: Some Thoughts on God, Reason, and Faith," in *Philosophers without Gods: Meditations on Atheism and the Secular Life*. Ed. Louise Antony. Oxford, Oxford University Press.

Gardiner, Georgi (2012). Understanding, Integration, and Epistemic Value. *Acta Analytica*. 27(2): 163-81.

Gardiner, Georgi (2015). "Teleologies and the Methodology of Epistemology," in *Epistemic Evaluation: Purposeful Epistemology*. Eds. David Henderson and John Greco. Oxford, Oxford University Press.

Geach, Peter (1956). Good and Evil. *Analysis*. 17(2): 33-42.

Gendler, Tamar Szabó (2002). Personal Identity and Thought-Experiments. *The Philosophical Quarterly*. 52(206): 34-54.

Gettier, Edmund (1963). Is Justified True Belief Knowledge? *Analysis*. 23: 121-3.

Gibbons, John (2013). *The Norm of Belief*. Oxford, Oxford University Press.

Ginet, Carl (2001). "Deciding to Believe," in *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford, Oxford University Press.

Goldman, Alvin (1967). A Causal Theory of Knowing. *The Journal of Philosophy*. 64(12): 357-72.

Goldman, Alvin (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy*. 73: 771-91.

Goldman, Alvin (1979). "What is Justified Belief?" in *Justification and Knowledge*. Ed. G. Pappas. Dordrecht, Reidel.

Goldman, Alvin (1986). *Epistemology and Cognition*. Cambridge, Harvard University Press.

Goldman, Alvin (1999). *Knowledge in a Social World*. Oxford, Oxford University Press.

Goldman, Alvin (2001). "The Unity of the Epistemic Virtues," in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford, Oxford University Press.

Goldman, Alvin and Olsson, Erik (2009). "Reliabilism and the Value of Knowledge," in *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Goldman, Alvin (2015). Reliabilism, Veritism, and Epistemic Consequentialism. *Episteme*. 12(2): 131-43.

Good, Irving John (1967). On the Principle of Total Evidence. *The British Journal for the Philosophy of Science*. 17(4): 319-21.

Greco, John (2003). "Knowledge as Credit for True Belief," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski. Oxford, Clarendon Press.

Greco, John (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge, Cambridge University Press.

Grimm, Stephen (2008). Epistemic Goals and Epistemic Values. *Philosophy and Phenomenological Research*. 77(3): 725-44.

Grimm, Stephen (2009). "Epistemic Normativity," in *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Haack, Susan (2001). "'The Ethics of Belief' Reconsidered," in *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford, Oxford University Press.

Harman, Gilbert (1986). *Change in View: Principles of Reasoning*. Cambridge, MIT Press.

Harman, Gilbert (1995). "Rationality," in *Thinking: Invitation to Cognitive Science*, Vol. Three. Eds. Edward Smith and Daniel Osherson. Cambridge, MIT Press.

Hart, H. L. A. (2008). *Punishment and Responsibility: Essays in the Philosophy of Law*. Second Edition. Oxford, Oxford University Press.

Hattiangadi, Anandi (2010). The Love of Truth. *Studies in History and Philosophy of Science*. 41(4): 422-32.

Hazlett, Allan (2013). *A Luxury of the Understanding: On the Value of True Belief*. Oxford, Oxford University Press.

Hazlett, Allan (2014). "Expressivism and Convention Relativism," in *Naturalizing Epistemic Virtue*. Eds. Abrol Fairweather and Owen Flanagan. Cambridge, Cambridge University Press.

Hazlett, Allan (2016). What Does 'Epistemic' Mean? *Episteme*. 13(4): 539-47.

Hempel, Carl (1965). "Aspects of Scientific Explanation," in *Aspects of Scientific Explanation and other Essays in the Philosophy of Science*. New York, Free Press.

Henderson, David and Horgan, Terence (2015). "What's the Point?" in *Epistemic Evaluation: Purposeful Epistemology*. Eds. David Henderson and John Greco. Oxford, Oxford University Press.

Hieronymi, Pamela (2005). The Wrong Kind of Reason. *The Journal of Philosophy*. 102(9): 437-57.

Hill, Thomas, Jr. (1996). "Moral Dilemmas, Gaps, and Residues," in *Moral Dilemmas and Moral Theory*. Ed. H. E. Mason. Oxford, Oxford University Press.

Hintikka, Jaakko (1999). "Interrogative Logic as a General Theory of Reasoning," in *Inquiry as Inquiry: A Logic of Scientific Discovery*. Dordrecht, Kluwer Academic Publishers.

Horwich, Paul (2006). The Value of Truth. *Nous*. 40: 347-60.

Howard-Snyder, Frances (2006). "Cannot" Implies "Not Ought." *Philosophical Studies*. 130(2): 233-46.

Humberstone, I. L. (1992). Direction of Fit. *Mind*. 101(401): 59-83.

Ichikawa, Jonathan and Jarvis, Benjamin (2013). *The Rules of Thought*. Oxford, Oxford University Press.

Ichikawa, Jonathan Jenkins (2014). Justification is Potential Knowledge. *Canadian Journal of Philosophy*. 44(2): 184-206.

Jackson, Frank and Pargetter, Robert (1986). Oughts, Options, and Actualism. *The Philosophical Review*. 95(2): 233-55.

James, William (1956). "The Will to Believe" in *The Will to Believe and Other Essays in Popular Philosophy*. New York, Dover Publications.

Jones, Ward (1997). Why Do We Value Knowledge? *American Philosophical Quarterly*. 34(4): 423-39.

Kappel, Klemens (2010). "On Saying that Someone Knows: Themes from Craig," in *Social Epistemology*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Keller, Simon (2004). Friendship and Belief. *Philosophical Papers*. 33(3): 329-51.

Kelly, Thomas (2002). The Rationality of Belief and Some Other Propositional Attitudes. *Philosophical Studies*. 110(2): 163-96.

Kelly, Thomas (2003). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research*. 66(3): 612-40.

Kelly, Thomas (2007). Evidence and Normativity: Reply to Leite. *Philosophy and Phenomenological Research*. 75(2): 465-74.

Kelly, Thomas (2014). "Evidence," in *The Stanford Encyclopedia of Philosophy*. Ed. Edward Zalta. URL = <https://plato.stanford.edu/entries/evidence>.

Kelp, Christoph (2011). What's the Point of 'Knowledge' Anyway? *Episteme*. 8: 53-66.

Kelp, Christoph (forthcoming). "Knowledge First Virtue Epistemology," in *Knowledge First: Approaches in Epistemology and Mind*. Eds. Adam Carter, Emma Gordon, and Benjamin Jarvis. Oxford, Oxford University Press.

Kornblith, Hilary (1985). *Naturalizing Epistemology*. Cambridge, MIT Press.

Kornblith, Hilary (1993). Epistemic Normativity. *Synthese*. 94(3): 357-76.

Kornblith, Hilary (2008). "Knowledge Needs No Justification," in *Epistemology: New Essays*. Ed. Quentin Smith. Oxford, Oxford University Press.

Korsgaard, Christine (1996). *The Sources of Normativity*. Cambridge, Cambridge University Press.

Kvanvig, Jonathan (2003a). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge, Cambridge University Press.

Kvanvig, Jonathan (2003b). Propositionalism and the Perspectival Character of Justification. *American Philosophical Quarterly*. 40(1): 3-18.

Kvanvig, Jonathan (2014a). "Truth is not the Primary Epistemic Goal," in *Contemporary Debates in Epistemology, Second Edition*. Eds. Matthias Steup, John Turri, and Ernest Sosa. West Sussex, Blackwell Publishing.

Kvanvig, Jonathan (2014b). "Epistemic Normativity," in *Epistemic Norms: New Essays on Action, Belief, and Assertion*. Eds. Clayton Littlejohn and John Turri. Oxford, Oxford University Press.

Lackey, Jennifer (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford, Oxford University Press.

Latus, Andrew (2000). Our Epistemic Goal. *Australasian Journal of Philosophy*. 78(1): 28-39.

Leary, Stephanie (2016). In Defense of Practical Reasons for Belief. *Australasian Journal of Philosophy*. 1-14.

Leite, Adam (2007). Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly's 'Epistemic Rationality as Instrumental Rationality: A Critique.' *Philosophy and Phenomenological Research*. 75(2): 456-64.

Lewis, David (1994). "Lewis, David: Reduction of Mind," in *A Companion to the Philosophy of Mind*. Ed. Samuel Guttenplan. Oxford, Blackwell Publishing.

Lipton, Peter (2004). *Inference to the Best Explanation*. Second Edition. London, Routledge Publishing.

Littlejohn, Clayton (2011). "Reasons and Belief's Justification," in *Reasons for Belief*. Eds. Andrew Reisner and Asbjørn Steglich-Petersen. Cambridge, Cambridge University Press.

Littlejohn, Clayton (2012). *Justification and the Truth-Connection*. Cambridge, Cambridge University Press.

Lynch, Michael (2004). *True to Life: Why Truth Matters*. Cambridge, MIT Press.

Lynch, Michael (2009). "The Values of Truth and the Truth of Values," in *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. London, Penguin Books.

Maguire, Barry (2016). The Value-Based Theory of Reasons. *Ergo*. 3(9): 233-62.

Markovits, Julia (2014). *Moral Reason*. Oxford, Oxford University Press.

McCain, Kevin (2014). *Evidentialism and Epistemic Justification*. New York, Routledge Publishing.

McCormick, Miriam (2015). *Believing Against the Evidence: Agency and the Ethics of Belief*. New York, Routledge Publishing.

McGrath, Matthew (2007). Memory and Epistemic Conservatism. *Synthese*. 157: 1-24.

McGrath, Matthew and Fantl, Jeremy (2013). "Truth and Epistemology," in *Virtuous Thoughts: The Philosophy of Ernest Sosa*. Ed. John Turri. Dordrecht, Springer.

McHugh, Conor (2011). What Do We Aim at When We Believe? *Dialectica*. 65(3): 369-92.

McHugh, Conor (2012). Belief and Aims. *Philosophical Studies*. 160(3): 425-39.

McKay, Ryan and Dennett, Daniel (2009). The Evolution of Misbelief. *Behavioral and Brain Sciences*. 32(6): 493-561.

Millar, Alan (2010). "Knowledge and Recognition," in *The Nature and Value of Knowledge: Three Investigations*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Miracchi, Lisa (2015). Competence to Know. *Philosophical Studies*. 172(1): 29-56.

Monmarquet, James (1986). The Voluntariness of Belief. *Analysis*. 46(1): 49-53.

Montefiore, Alan (1958). "Ought" and "Can." *Philosophical Quarterly*. 8(30): 24-40.

Moore, G. E. (1903). *Principia Ethica*. Cambridge, Cambridge University Press.

Moser, Paul (1985). *Empirical Justification*. Dordrecht, Reidel Publishing.

Nolfi, Kate (forthcoming). "Why Evidence (and Only Evidence) Can Justify Belief," in *Normativity: Epistemic and Practical*. Eds. Conor McHugh, Jonathan Way, and Daniel Whiting. Oxford, Oxford University Press.

Nottelmann, Nikolaj (2008). "The Present and Future State of Epistemic Deontologism," in *New Waves in Epistemology*. Eds. Vincent Hendricks and Duncan Pritchard. London, Palgrave Macmillan.

Nozick, Robert (1993). *The Nature of Rationality*. Princeton, Princeton University Press.

Oddie, Graham (1997). Conditionalization, Cogency, and Cognitive Value. *The British Journal for the Philosophy of Science*. 48(4): 533-41.

Parfit, Derek (1984). *Reasons and Persons*. Oxford, Clarendon Press.

Parfit, Derek (2011). *On What Matters, Volume One*. Oxford, Oxford University Press.

Pascal, Blaise (1966). *Pensees*. Oxford, Penguin Publishing.

Pereboom, Derk, ed., (2009). *Free Will*. Indianapolis, Hackett Publishing.

Piller, Christian (2006). Content-Related and Attitude-Related Reasons for Preferences. *Royal Institute of Philosophy Supplement*. 81(59): 155-82.

Piller, Christian (2009). Desiring the Truth and Nothing But the Truth. *Nous*. 43(2): 193-213.

Plantinga, Alvin (1993a). *Warrant and Proper Function*. Oxford, Oxford University Press.

Plantinga, Alvin (1993b). *Warrant: The Current Debate*. Oxford, Oxford University Press.

Plato (1961). *The Collected Dialogues of Plato*. Eds. Edith Hamilton and Huntington Cairns. Trans. W. K. C. Guthrie. New York, Bollingen Foundation.

Plunkett, David (2015). Which Concepts Should We Use?: Metalinguistic Negotiations and the Methodology of Philosophy. *Inquiry*. 58(7): 828-74.

Preston-Roedder, Ryan (2013). Faith in Humanity. *Philosophy and Phenomenological Research*. 87(3): 664-87.

Pritchard, Duncan (2010). "Knowledge and Understanding," in *The Nature and Value of Knowledge: Three Investigations*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford, Oxford University Press.

Pritchard, Duncan (2014). "Truth as the Fundamental Epistemic Good," in *The Ethics of Belief*. Eds. Jonathan Matheson and Rico Vitz. Oxford, Oxford University Press.

Quine, W. V. O. (1986). "Reply to Morton White," in *The Philosophy of W.V. Quine*. Eds. Lewis Hahn and Paul Schilpp. La Salle, Open Court Publishing.

Rabinowicz, Wlodek and Roennow-Rasmussen, Toni (1999). A Distinction in Value: Intrinsic and for Its Own Sake. *Proceedings of the Aristotelian Society*. 100: 33-49.

Rabinowicz, Wlodek and Roennow-Rasmussen, Toni (2003). Tropic of Value. *Philosophy and Phenomenological Research*. 66: 389-403.

Radford, Colin (1966). Knowledge—By Examples. *Analysis*. 27: 1-11.

Railton, Peter (1994). Truth, Reason, and the Regulation of Belief. *Philosophical Issues*. 5: 71-93.

Rawls, John (1963). Constitutional Liberty and the Concept of Justice. *Nomos*. 6: 98-125.

Raz, Joseph (1999). *Practical Reason and Norms*. Oxford, Oxford University Press.

Reisner, Andrew (2009). The Possibility of Practical Reasons for Belief and the Wrong Kind of Reasons Problem. *Philosophical Studies*. 145(2): 257-72.

Reisner, Andrew (forthcoming). "Pragmatic Reasons for Belief," in *The Oxford Handbook of Reasons and Normativity*. Ed. Daniel Star. Oxford, Oxford University Press.

Rescher, Nicholas (1984). *Pascal's Wager*. South Bend, University of Notre Dame Press.

Rescorla, Michael (2009). Cognitive Maps and the Language of Thought. *British Journal for the Philosophy of Science*. 60(2): 377-407.

Ridge, Michael (2013). Getting Lost on the Road to Larissa. *Nous*. 47(1): 181-201.

Riggs, Wayne (2002). Reliability and the Value of Knowledge. *Philosophy and Phenomenological Research*. 64(1): 79-96.

Riggs, Wayne (2003). Balancing Our Epistemic Goals. *Nous*. 37(2): 342-52.

Riggs, Wayne (2008). "The Value Turn in Epistemology," in *New Waves in Epistemology*. Eds. Vincent Hendricks and Duncan Pritchard. London, Palgrave Macmillan.

Rinard, Susanna (2015). Against the New Evidentialists. *Philosophical Issues*. 25(1): 208-23.

Rinard, Susanna (2017). No Exception for Belief. *Philosophy and Phenomenological Research*. 94(1): 121-43.

Russell, Bruce (2001). "Epistemic and Moral Duty," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Ed. Matthias Steup. Oxford, Oxford University Press.

Ryan, Sharon (2003). Doxastic Compatibilism and the Ethics of Belief. *Philosophical Studies*. 114(1): 47-79.

Sanchez, Richard (2013). Belief in a Meaningless Multiverse. *British Journal of Behavioral Sciences*. 13(2): 67-83.

Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, Harvard University Press.

Schroeder, Mark (2007). *Slaves of the Passions*. Oxford, Oxford University Press.

Schroeder, Mark (2012). The Ubiquity of State-Given Reasons. *Ethics*. 122(3): 457-88.

Schroeder, Mark (forthcoming). "The Unity of Reasons," in *The Oxford Handbook of Reasons and Normativity*. Ed. Daniel Star. Oxford, Oxford University Press.

Schroeder, Timothy (2003). Donald Davidson's Theory of Mind is Non-Normative. *Philosophers' Imprint*. 3(1): 1-14.

Searle, John (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge, Cambridge University Press.

Setiya, Kieran (2013). Epistemic Agency: Some Doubts. *Philosophical Issues*. 23(1): 179-98.

Shah, Nishi (2003). How Truth Governs Belief. *The Philosophical Review*. 112(4): 447-82.

Shah, Nishi and Velleman, J. David (2005). Doxastic Deliberation. *The Philosophical Review*. 114(4): 497-534.

Shah, Nishi (2006). A New Argument for Evidentialism. *The Philosophical Quarterly*. 56(225): 481-98.

Siegel, Harvey (2005). Truth, Thinking, Testimony and Trust: Alvin Goldman on Epistemology and Education. *Philosophy and Phenomenological Research*. 71(2): 345-66.

Skorupski, John (2010). *The Domain of Reasons*. Oxford, Oxford University Press.

Skyrms, Brian (1990). *The Dynamics of Rational Deliberation*. Cambridge, Harvard University Press.

Slote, Michael (1989). *Beyond Optimizing: A Study of Rational Choice*. Cambridge, Harvard University Press.

Slote, Michael (1990). Ethics without Free Will. *Social Theory and Practice*. 16(3): 369-83.

Smithies, Declan (2015). "Why Justification Matters," in *Epistemic Evaluation: Purposeful Epistemology*. Eds. David Henderson and John Greco. Oxford, Oxford University Press.

Solomon, David (2003). "Virtue Ethics: Radical or Routine?" in *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski. Oxford, Oxford University Press.

Sosa, Ernest (2001). "For the Love of Truth?" in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford, Oxford University Press.

Sosa, Ernest (2003). "The Place of Truth in Epistemology," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski. Oxford, Oxford University Press.

Sosa, Ernest (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford, Oxford University Press.

Sosa, Ernest (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. Oxford, Oxford University Press.

Sosa, Ernest (2015). *Judgment and Agency*. Oxford, Oxford University Press.

Steglich-Petersen, Asbjørn (2006). No Norm Needed: On the Aim of Belief. *The Philosophical Quarterly*. 56(225): 499-516.

Steglich-Petersen, Asbjørn (2009). Weighing the Aim of Belief. *Philosophical Studies*. 145(3): 395-405.

Steglich-Petersen, Asbjørn (2011). "How to be a Teleologist about Epistemic Reasons," in *Reasons for Belief*. Eds. Andrew Reisner and Asbjørn Steglich-Petersen. Cambridge, Cambridge University Press.

Steglich-Petersen, Asbjørn (2013). "Truth as the Aim of Epistemic Justification," in *The Aim of Belief*. Ed. Timothy Chan. Oxford, Oxford University Press.

Steup, Matthias (2000). Doxastic Voluntarism and Epistemic Deontology. *Acta Analytica*. 15(1): 25-56.

Steup, Matthias (2008). Doxastic Freedom. *Synthese*. 161(3): 375-92.

Stroud, Sarah (2006). Epistemic Partiality in Friendship. *Ethics*. 116(3): 498-524.

Swain, Marshall (1981). *Reasons and Knowledge*. Ithaca, Cornell University Press.

Swinburne, Richard (1999). *Providence and the Problem of Evil*. Oxford, Oxford University Press.

Swinburne, Richard (2001). *Epistemic Justification*. Oxford, Clarendon Press.

Sylvan, Kurt (forthcoming). Veritism Unswamped. *Mind*.

Taylor, S. E. and Brown, J. D. (1988). Illusion and Well-Being: A Social Psychological Perspective on Mental Health. *Psychological Bulletin*. 103(2): 193-210.

Thomson, Judith Jarvis (1976). Killing, Letting Die, and the Trolley Problem. *Monist*. 59: 204-17.

Thomson, Judith Jarvis (2001). *Goodness and Advice*. Ed. Amy Gutmann. Princeton, Princeton University Press.

Thomson, Judith Jarvis (2008). *Normativity*. Chicago, Open Court Publishing.

Treanor, Nick (2014). Trivial Truths and the Aim of Inquiry. *Philosophy and Phenomenological Research*. 89(3): 552-9.

Turri, John (2010). On the Relationship between Propositional and Doxastic Justification. *Philosophy and Phenomenological Research*. 80(2): 312-26.

Velleman, J. David (2000). *The Possibility of Practical Reason*. Oxford, Oxford University Press.

Weatherson, Brian (2008). Deontology and Descartes' Demon. *The Journal of Philosophy*. 105(9): 540-69.

Wedgwood, Ralph (2002). The Aim of Belief. *Nous*. 36: 267-97.

Wedgwood, Ralph (2013a). "The Right Thing to Believe," in *The Aim of Belief*. Ed Timothy Chan. Oxford, Oxford University Press.

Wedgwood, Ralph (2013b). Rational "Ought" Implies "Can." *Philosophical Issues*. 23(1): 70-92.

Whitcomb, Dennis (2007). *An Epistemic Value Theory*. PhD dissertation, Rutgers University. Unpublished manuscript. Accessed at Rutgers University Community Repository.

Whitcomb, Dennis (2014). "Can there be a Knowledge-First Ethics of Belief?" in *The Ethics of Belief*. Eds. Jonathan Matheson and Rico Vitz. Oxford, Oxford University Press.

Whiting, Daniel (2010). Should I Believe the Truth? *Dialectica*. 64(2): 213-24.

Whiting, Daniel (2013). "Nothing but the Truth: On the Norms and Aims of Belief," in *The Aim of Belief*. Ed. Timothy Chan. Oxford, Oxford University Press.

Whiting, Daniel (2014). "Reasons for Belief, Reasons for Action, the Aim of Belief, and the Aim of Action," in *Epistemic Norms: New Essays on Action, Belief, and Assertion*. Eds. Clayton Littlejohn and John Turri. Oxford, Oxford University Press.

Williams, Bernard (1973). "Deciding to Believe," in *Problems of the Self*. Cambridge, Cambridge University Press.

Williams, Bernard (1981). "Internal and External Reasons," in *Moral Luck: Philosophical Papers, 1973-1980*. Cambridge, Cambridge University Press.

Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford, Oxford University Press.

Wood, Allen (2008). The Duty to Believe According to the Evidence. *International Journal for Philosophy of Religion*. 63(1): 7-24.

Zagzebski, Linda (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge, Cambridge University Press.

Zagzebski, Linda (2000). "From Reliabilism to Virtue Epistemology," in *Knowledge, Belief, and Character*. Ed. Guy Axtell. Lanham, Rowman and Littlefield.

Zagzebski, Linda (2001). "Recovering Understanding," in *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford, Oxford University Press.

Zagzebski, Linda (2003). The Search for the Source of Epistemic Good. *Metaphilosophy*. 34: 12-28.

Zimmerman, Michael (1996). *The Concept of Moral Obligation*. New York, Cambridge University Press.