

Copyright

by

Brian William Pollex

2018

**The Report Committee for Brian William Pollex
Certifies that this is the approved version of the following Report:**

Locating Epistemic Value

**APPROVED BY
SUPERVISING COMMITTEE:**

Sinan Dogramaci, Supervisor

David Sosa
Ray Buchanan
Matt Evans
Cory Juhl

Locating Epistemic Value

by

Brian William Pollex

Report

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Master of Arts

The University of Texas at Austin

December 2018

Dedication

For C.S.

Acknowledgements

I am deeply grateful for the patience and support of my committee: Sinan Dogramaci, supervisor, David Sosa, reader, Ray Buchanan, Matthew Evans, and Cory Juhl. I would also like to thank my many other teachers and colleagues in philosophy, without whose warmth, insight, and intelligence the material that follows would suffer many more defects than it assuredly already does. My deepest thanks and gratitude go to my family and friends, without whose support I would not have been able to produce this report.

Abstract

Locating Epistemic Value

Brian William Pollex, M.A.

The University of Texas at Austin, 2018

Supervisor: Sinan Dogramaci

Abstract: Many epistemologists are attracted to the idea that knowledge is valuable in a way that stands out when compared to the value of other doxastic attitudes. Some philosophers, including Linda Zagzebski, Jonathan Kvanvig, Richard Swinburne, and Duncan Pritchard have objected to contemporary theories of knowledge on the grounds that the analyses these theories offer fail to sufficiently distinguish knowledge from mere true belief. One particularly clear instance of this is the Swamping Problem for Process Reliabilism. Goldman and Olsson try to respond to the Swamping Problem on behalf of Reliabilism. In what follows, I first try to motivate and defend an approach to epistemic value. Then I review the Swamping Problem and evaluate Goldman and Olsson's response to it. Finding their response unsatisfying, I suggest that we try to satisfy Zagzebski et al. by introducing a theory of understanding which answers to the intuitions many endorse about epistemic value.

Table of Contents

List of Tables	ix
INTRODUCTION.....	1
I.1 Kvanvig and Goldman & Olsson.....	1
I.2 Structure and Aims	3
LOCATING EPISTEMIC VALUE	8
Chapter 1: How to Think About Epistemic Value.....	8
1.1 The Guiding Intuition	8
1.2 Why Accept Something Like the Guiding Intuition?	13
1.3 Pritchard’s Three Value Problems	20
1.4 A Dispute Between Frameworks	24
Chapter 2: Truth, Belief, Knowledge, and Value	28
2.1 A Disclaimer About the Inquiry into “Mere True Belief”	28
2.2 Knowledge and Instrumental Value.....	31
2.3 Acting Without Knowledge (Or Belief)	35
2.4 Mere True Belief, Veritism, and Epistemic “Junk”	38
Chapter 3: The Swamping Problem and the Conditional Probability Solution	40
3.1 The Swamping Problem.....	40
3.2 The Conditional Probability Solution	46
3.3 Evaluating the Success of the Conditional Probability Solution	48
3.4 Framing the Generality Problem.....	52
3.5 The Conditional Probability Solution and the Generality Problem	56

Chapter 4: Understanding	62
4.1 Kinds of Understanding	62
4.2 On Objectual Understanding.....	65
References.....	72

List of Tables

Table 1: Theoretical approaches to epistemic understanding	65
--	----

INTRODUCTION

I.1 Kvanvig and Goldman & Olsson

Much of mainstream contemporary epistemology has been occupied with crafting a satisfactory account of knowledge. Comparatively little has been said about what it is that makes knowledge valuable. This is somewhat surprising given the overwhelming volume of work that has been produced giving increasingly complex counterexamples, revisions, and re-revisions of competitor theories of knowledge.

Although epistemic value is an appealing topic, there is an obvious obstacle in the way of serious investigation. How can we productively inquire into the value of knowledge without anything even approximating a consensus on the correct theory of knowledge, or even the kind of value knowledge would have if the correct theory were in place?¹ In *The Value of Knowledge and the Pursuit of Understanding*, Jonathan Kvanvig offers an implicit methodological suggestion by proceeding piecemeal, discussing several serious contender theories in turn. My project is in part an endorsement of this approach.

Kvanvig considers whether or not Process Reliabilism (hereafter sometimes simply Reliabilism) is compatible with a plausible account of the value of knowledge. He uses Reliabilism as a test case for Externalist theories of knowledge, and suggests that conclusions he draws about Reliabilism can be extended (with minor alterations) to other species of the Externalist genus. Kvanvig argues that no satisfying account of epistemic value is forthcoming from the Reliabilist, and so suggests that prospects look dim for Externalism more generally. He also suggests that certain Internalist theories might have

¹ Kvanvig and others have defended the view that knowledge has “categorical” value (Kvanvig 2003) while prefer alternative accounts of the nature of knowledge’s value. See Section 1.1 for more discussion.

access to more promising theoretical tools for explaining epistemic value. He concludes with a tentative investigation of whether Understanding, a different epistemic state than knowledge, might fare better than knowledge as the bearer of epistemic value.²

So far as the Externalist, and specifically the Reliabilist, is concerned, Kvanvig's picture looks like this: on Externalist theories of knowledge, like Reliabilism, knowledge is not plausibly thought of as distinctively epistemically valuable.³ But an Externalist theory of *knowledge* (like Reliabilism) might be paired with or supplemented with an account of another kind of epistemic state, for example, Understanding, which might be the actual bearer of the distinctive epistemic value we had hoped to locate in knowledge.

In order to push this line, Kvanvig relies on an argument presented by Linda Zagzebski and several others usually called the Swamping problem.⁴ Kvanvig's agenda in his book is broad, but his treatment of the Swamping problem in particular has garnered response. In their 2009 paper *Reliabilism and the Value of Knowledge*, Alvin Goldman and Erik J. Olsson respond on behalf of Reliabilism to the Swamping problem as presented by Kvanvig and others. Goldman and Olsson attempt to show that adopting a Process Reliabilist theory of knowledge is compatible with holding that knowledge has a special value, albeit of a different kind than Kvanvig is after. Goldman and Olsson's piece itself

² Depending on how one conceives of the task at hand, this may seem very unsatisfying; if one takes the issue to be the narrow matter of discovering what it is that makes *knowledge* valuable, then of course investigation of any other epistemic state is a waste of time. But there is a broader way to understand interest in epistemic value. See Section 1.2 for discussion of the move from the *Guiding Intuition* to the *Weak Guiding Intuition*.

³ More on what it means to be "distinctively epistemically valuable" will follow in Section 1.1.

⁴ Swinburne and Zagzebski each discuss related earlier versions of the problem, Swinburne in his 1998 book *Providence and the Problem of Evil* and Zagzebski in her 1996 book *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Ward Jones is credited with a version of the problem in 1997. By Kvanvig 2003 the problem was well known, and Goldman and Olsson 2009 compare several versions of the swamping problem.

has spawned further discussion about the prospects for Reliabilism with regard to epistemic value.⁵

I.2 Structure and Aims

My work here will be broken up into four main sections. In chapters one and two I will explain and defend my approach to the literature on epistemic value. After this, I have two principal aims. In chapter three, I will show that Goldman and Olsson's Conditional Probability Solution does not satisfactorily respond to the Swamping Problem, leaving Process Reliabilism facing a serious issue with regards to epistemic value. In chapter four, I will propose an account of Understanding which I take to be well-suited to a partnered role for theories of knowledge like Reliabilism.

As mentioned above, my project is methodologically in the spirit of Kvanvig's, although my own views are in some ways less radical than his and in other ways more. Unlike Kvanvig, I will constrain my remarks to Reliabilism directly rather than attempting to extend them to cover all possible Externalist approaches. In this respect my goals are more modest. However my views about the role of Understanding for the Reliabilist are somewhat more strident than Kvanvig's. Rather than just raising the possibility that certain Externalist theories of knowledge might be supplemented by an epistemic value-bearing account of Understanding, I will offer a more detailed, substantial proposal for the kind of epistemic state that would be best suited to play this partnered role for Reliabilism.

Before beginning in earnest, I should note that Process Reliabilism does face serious and well-known extensional objections, with pressure coming from both directions. Cases like BonJour's Norman suggest to some that reliability does not suffice for

⁵ See the exchanges between Olsson, Jäger, and others in *Theoria*.

justification, and proponents of Demon-World style objections oppose the necessity of reliability for justification. There are also other pressing issues for Reliabilism simply *qua* theory of knowledge, like the Generality problem.

That said, I will not be directly addressing objections to Process Reliabilism *qua* theory of knowledge, as I am not attempting to motivate the acceptance of Reliabilism. I will discuss the Generality Problem in section 3.4, but only because of a connection to the debate over epistemic value. Relatedly, though many contributors⁶ to the literature note that a Gettier condition is needed in order to give the most plausible version of a Reliability theory of knowledge, in the context of the debate over epistemic value most (though not all) have preferred to focus on a simplified version of the view that omits this condition in order to focus more narrowly on the relevant aspects of the theory.⁷ I will do likewise.

These considerations raise an important question. Why focus on a simplified version of a beleaguered view? Like Kvanvig, I think an Externalist approach to the theory of knowledge is well-supported on grounds independent of epistemic value, and that because of this, serious consideration of Reliabilism--even simplified Reliabilism--is a promising way to make progress with developing issues in the debate over epistemic value. One strong reason to favor an externalist approach in the theory of knowledge has to do with dialectical pressure to get certain cases right: non-inferential knowledge, perceptual knowledge, and other cases involving non-reflecting or unsophisticated agents suggest that any account requiring cognitively sophisticated conditions be met in order for one to know, is likely to be highly revisionary for that very reason. An externalist account that merely

⁶ Goldman and Olsson, Kvanvig, Jager, Zagzebski, etc.

⁷ Goldman and Olsson write that, "...the idea of knowledge depending on the existence of a reliable connection is the central one behind reliabilism, and it would be unfortunate for the theory if that very component failed to produce an added value" (2009, pg. 155). For a contrasting approach, Riggs holds that the Gettier condition may be a feature of interest to philosophers hoping to make progress on questions of epistemic value. Georgi Gardiner (2017) argues that the modal features of some accounts could not explain knowledge's value.

constrains belief formation or the modal profile of a given belief seems more likely to capture these important cases. And Reliabilism, though it faces difficulties, is one such view.

Whether one prefers a sophisticated Reliabilist account of knowledge, such as the Agent- or Virtue-Reliabilist theories, some other Externalist account of knowledge, an Internalist account, or (perhaps like Bonjour, Alston, or on my reading, Klein) is skeptical of unified accounts of knowledge,⁸ the debate that has developed around Reliabilism is instructive. It offers a relatively clear opportunity to see how several major figures think about the connection between knowledge, true belief, and epistemic value.

In chapter four I will turn to the positive part of my project: presenting an account of Understanding that I regard as distinctively epistemically valuable. My hope is that this suggestion will be made more intriguing than it would be were it totally free-standing by offering it as an extension of a familiar (if not widely endorsed) theory of knowledge, and an extension which ameliorates a serious concern for that view.

Unfortunately, like knowledge, the proper analysis of Understanding is hotly debated. On some prominent views, Understanding is a special kind of knowledge--but only in the sense of having a specific topic, and not special in the sense of being more difficult to acquire than knowledge of other topics. On these views, very roughly, understanding is knowledge of causes.⁹ I take it that this is meant to capture the intuitive idea that, at least in many cases, to understand something is to know *why* that thing is the case, or why that proposition is true.

⁸ See Alston's 2005 book *Beyond "Justification": Dimensions of Epistemic Evaluation* for more information about his view. Klein's infinitism is meant, on my view, primarily to capture what he calls "adult human knowledge". This leaves him without a clear response to "animal knowledge" cases. See his 1999 paper, "Human knowledge and the infinite regress of reasons".

⁹ See Kvanvig's 2003 discussion of propositional understanding; compare Elgin and Riggs' Appendices C and D to Haddock, Millar & Pritchard's 2009 collection *Epistemic Value*.

On other views understanding is a relationship between an agent and a domain, subject, topic, or craft (e.g., physics), rather than an individual proposition. On these views understanding in its epistemic guise is primarily related to the phenomenon inquired after or reported about with sentences like “does anyone here understand calculus?”, or “Hannah really understands contemporary U.S. politics”. Some (usually, but not only, members of the first camp) have argued that Understanding is factive while others (usually but not only members of the second camp) hold that it is best understood to be non-factive or “indirectly factive”.¹⁰

While I will be discussing a non-propositional, non-factive account of Understanding and will discuss some intuitions and examples that I take to motivate such an account, I do not intend to try to argue that this is the best or only plausible analysis of an epistemic state denoted by our use of the word ‘understanding’. This is because I take it that ‘understanding’ has familiar, entrenched, properly epistemic uses, some of which are propositional (e.g., “Claire finally *understood that* it is wrong to break promises without a good reason”) and others which are non-propositional or only propositional in an indirect sense (e.g., “Sarah really *understands* calculus”).¹¹

I do not advocate a non-propositional, non-factive account of understanding as the correct analysis of (epistemic) understanding *tout court*, but I do think that the state referred to by the non-propositional, non-factive uses of ‘understanding’ (ie., objectual understanding) is the best candidate for the epistemic state that bears an epistemic value distinctive from that borne by states like knowledge and mere true belief.¹²

¹⁰ See Zagzebski (2001).

¹¹ That is, indirectly propositional by virtue of being related to propositions indirectly, while being directly related to e.g. a group of propositions. A state that is indirectly propositional in this sense is propositional by virtue of being (indirectly) related to propositions, but fails to be directly propositional by directly relating agents to an object that is not itself a proposition (generally, a body of propositions).

¹² See also section 2.1.

I hope to lessen the potential impact of some of these complications by presenting the non-factive, non-propositional account as a doxastic state especially well-suited to complement a Reliabilist (or other externalist) theory of knowledge when it comes to questions of epistemic value. Having argued that certain analyses of knowledge (specifically Reliabilist ones) fail to satisfy some forceful intuitions about epistemic value, my proposal is not to abandon an independently well-motivated theory of knowledge just because that theory fits poorly with some desiderata, but instead to move to a more ecumenical epistemology by complementing a well-supported theory of knowledge with an account of understanding, where the latter *does* satisfy (a version of) the value-theoretic desiderata that our otherwise well-supported theory of knowledge fails to vindicate. In this way we can hopefully both have and eat our cake.

In short, my proposal is to respond to the difficulties presented by questions of epistemic value in a positive, rather than negative, way. Rather than wielding intuitions about value as objections to popular, well-supported theories of knowledge, I think that our interest in (and familiarity with) epistemic value ought to guide our epistemic theorizing, even if this takes us some distance away from familiar questions about knowledge.

LOCATING EPISTEMIC VALUE

Chapter 1: How to Think About Epistemic Value

1.1 The Guiding Intuition

Suppose, for the sake of expository simplicity and harmony with the existing literature, that there are such things as propositions roughly as many philosophers take there to be, and that the sorts of things they get up to are the sorts of things many hold that they do: they are abstracta, can feature in thought, are the objects of our beliefs, can be expressed by utterances, and are the proper bearers of truth and falsity. In addition to belief, they are also the objects to which we bear other propositional attitudes. For some proposition P, we might hope that P, fear that P, withhold from P, believe that P, disbelieve or reject that P, accept or suppose that P,¹³ predict that P, and so on.¹⁴ More contentiously, we might *understand* that P.¹⁵ Leaving aside understanding, each of belief, disbelief or rejection, acceptance or supposition, and prediction are examples of *doxastic* propositional attitudes, a term I understand to mean belief-like: related to an epistemic commitment, in one sense or another, that the world is thus-and-so, is or is not or might be as P says it is.

It is crucial to following this debate to notice at the outset that the relevant literature is organized around an investigation into doxastic propositional attitudes. To the extent that

¹³ See Patrick (1990) for discussion of Van Fraassen's notion of acceptance and the debate over its relation to belief.

¹⁴ There are several substantive debates about the nature of propositions and belief that raise deep and interesting questions about, for example, whether or not propositions figure in thought, or whether beliefs should be thought of as some form of inner assent, a way of taking the world to be thus-and-so, or instead as mere dispositions to act in certain ways. Largely because of the developing nature of the literature on epistemic value, these questions and others seem to have been set aside in order to try to make progress in laying out the fundamental issues at stake in this debate. This does not mean that these questions do not bear on the inquiry into epistemic value; on the contrary, they represent open avenues for future research. Generally though I will keep to this simplified picture for ease of exposition and consonance with the existing literature.

¹⁵ The notion of understanding in play here and in the remainder of the piece is epistemic rather than linguistic.

there is contemporary discussion of the older (ancient) picture of epistemic value or virtue, on which *wisdom* or some other character trait is the chief epistemic virtue, this discussion is largely separate from the investigation I am joining into doxastic propositional attitudes. I'd like to remain neutral on the issue of epistemic virtue in this sense in this piece. For my own part, I take it to be an obvious truth that some people are wise and some foolish, and that this judgment is at least as clearly evaluative as it is epistemic. It seems very plausible that a complete picture of what it is to do well epistemically must include some explanation of this and perhaps other epistemic virtues. I feel similarly about the view that rationality holds some special place in the epistemic domain. My aim here is not to furnish a complete picture of epistemic value. I merely wish to clarify the nature of the value borne by different doxastic propositional attitudes--specifically, to propose a shift in our attention away from knowledge.

Many epistemologists are attracted to the idea that, among the doxastic propositional attitudes, knowledge is *valuable*. Some would go even further, claiming that knowledge is not just valuable, but *distinctively* epistemically valuable—valuable in a way not shared by other doxastic states. Due to the important motivational role this intuition seems to play in the literature on epistemic value, I'll call this the *Guiding Intuition*.

Guiding Intuition: Knowledge is a distinctively epistemically valuable state.

This formulation is meant to capture the sense many epistemologists seem to have, which some explicitly state, that *knowledge* is preeminent in terms of value in the epistemic domain, while also being acceptable to philosophers (for example, epistemic virtue theorists) who prefer a more inclusive overall picture of the epistemic with room for, and indeed high esteem for, states like wisdom. So long as such theorists conceive of

knowledge as having a value that is unique among the doxastic attitudes and for that reason distinctive from the value that inheres in other features (e.g., wisdom, rationality), such theorists should also be able to accept the Guiding Intuition.

While I take it to be valuable to identify a common ground shared between such diverse epistemological approaches, if agreement is limited to the Guiding Intuition, then that agreement will not be particularly informative. The specific sense of value at stake in the Guiding Intuition remains unclear even taking for granted a traditional theory of knowledge.¹⁶ Without even a sketch of the kind of value one imputes to knowledge, it is unclear how to determine whether this value is actually distinctive or not. One extremely modest thesis would be that at least some knowledge is sometimes *instrumentally* or *extrinsically* valuable for certain people, given that those people have certain (perhaps very common) psychological dispositions and practical aims. A very ambitious alternative that might also be defended would be the view that all knowledge is *finally* or *intrinsically* valuable. Either of these proposals, if combined with a reason to think that this value is distinct from that associated with other doxastic attitudes, would vindicate the *Guiding Intuition*.

To be clear, vindicating the Guiding Intuition in this way is not the approach of all disputants in the literature. Kvanvig, for example, conceives of our task as finding out what makes knowledge “categorically” valuable: that is, valuable just in virtue of being knowledge. And surely a class of epistemic states could have a value that is distinctive without this value being categorical in Kvanvig’s sense. (For example, if only knowledge had some special epistemic value, but not *all* cases of knowledge had this feature.) It seems clear that the categorical value Kvanvig has in mind will also be distinctive in my sense.

¹⁶ By this I mean taking for granted a belief-first, rather than knowledge-first, analysis. More discussion of the implications of a knowledge-first framework will be discussed in section 2.1.

So it is clear that Kvanvig is not quite animated by the Guiding Intuition, although he does seem motivated by a deeply related intuition, one that seems even more demanding than the Guiding Intuition.

For their part, Goldman and Olsson suggest instead that we should aim to show that knowledge is valuable in what they call a “generic” sense, admitting of some (but not many) counterexamples.¹⁷ While they strive to show that (reliabilist) knowledge is more valuable than mere true belief, they do not specifically claim that the value of this state is distinctive, although it seems to me that Olsson at least would be friendly to the suggestion.¹⁸ So while neither party has explicitly endorsed the Guiding Intuition, it seems fair to characterize it as neutral territory in this dispute.

There is a challenge to Kvanvig’s categorical approach related to so-called “junk beliefs” or “junk knowledge”, the inspiration for which is usually attributed to Harman.¹⁹ If one takes this or other related problems seriously one might move away from the idea that knowledge is categorically valuable and instead hold that only some kinds or classes of knowledge (the non-junky ones) are valuable. (Such a view could be similar to the “generic” view endorsed by Goldman and Olsson.) A respect for the Harmanian “junk” problem combined with the conviction that knowledge is special within the epistemic domain is the straightforward reason to frame the Guiding Intuition as I have. Further, since the categorical value interpretation is controversial and since I mean for the Guiding

¹⁷ Their support for this claim consists of an appeal to dispositions to assent to natural language claims by ordinary users. It is worth noting that, as Goldman and Olsson are well aware, their approach will not deliver the stronger result, sought after by Kvanvig and others, that knowledge is categorically valuable. For a challenge related to the categorical value view, see discussion of the problem posed by so-called “junk beliefs” in section 2.4.

¹⁸ Goldman’s preferred response in the 2009 Goldman and Olsson paper (called “Value Autonomization”) to epistemic value claims is to provide a sort of epistemic error theory, and so it does not seem likely that he would endorse the Guiding Intuition.

¹⁹ Though Harman’s discussion in his 1986 book *Change in View* helped frame the debate, many others have contributed since then including important recent work, e.g. Jane Friedman (2018).

Intuition to have broad appeal, I prefer not burden the Guiding Intuition with the categorical value claim.

I introduced the *Guiding Intuition* because I think the most useful way to conceive of the literature on the value of knowledge is as an attempt to precisify the *Guiding Intuition* by discovering what kinds of value inhere in knowledge. Some distinctions will better equip us for our search. We will need the familiar distinction between *intrinsic* (or *final*) value and *extrinsic* (or *relational*) value.²⁰ I will follow Michael Zimmerman in characterizing intrinsic value as value some thing or state has “in itself,” “as such,” “for its own sake,” or “in its own right”, and characterizing extrinsic value as “value that is not intrinsic”. This way of dividing up the terrain, though exhaustive, is not particularly informative. We can say more about extrinsic value.

Instrumental value is one familiar kind of extrinsic value. Roughly, something has instrumental value given that it is a means or instrument to attaining something else of value. For example, a ticket to a baseball game might have instrumental value for me if it was a convenient means toward my goal of watching the game. In this example, I do not value the ticket in itself, nor is it intrinsically valuable, but it is instrumentally valuable as an instrument or means to something else I aim for.

We will also need recourse to the notion of *indicator* value, a second class of extrinsic value. Something, for example, medical test results, have indicator value when they indicate the presence of something of intrinsic value, for example, health. The test results are not a means to health, but they signify its presence. These distinctions in hand, we can return shortly to the *Guiding Intuition* and our attempt to precisify it. First, I’d like

²⁰ Although these two pairs of terms generally go together, there are edge cases in which, for example, something is finally valued for its extrinsic or relational properties, such as the familiar case of Princess Diana’s dress. See Korsgaard (1983).

to gesture to the motivations for proposals like the *Guiding Intuition*, and then briefly address a different approach to epistemic value.

1.2 Why Accept Something Like the Guiding Intuition?

Presumably most of us already accept that there are some epistemic states or doxastic attitudes that bear instrumental value.²¹ Why should we go further? That is, why should we accept that which the Guiding Intuition presupposes--that there is some value in the epistemic domain that is important and distinctive from the instrumental value that is typical of (or at least possibly borne by) many different kinds of doxastic attitudes?²² One response, offered by Ernest Sosa and endorsed by Kurt Sylvan,²³ requires the adoption of what Sosa calls “the epistemic point of view”. The epistemic point of view is a perspective on evaluative questions that treats the epistemic domain as a sphere of evaluative inquiry unto itself, and leaves aside the question of whether or not those success and failure conditions ought to be taken to be weighty, in the normative or evaluative sense, in a domain-independent way.

We are already intimately familiar with the epistemic point of view. It is this perspective we adopt when we negatively evaluate strategies like random guesswork and lazy adherence to authority, as well as states like simultaneously holding badly incoherent beliefs. Likewise for our positive evaluations of knowledge, wisdom, rationality, the ability

²¹ I don't see why one would deny this so long as they accept that there are a reasonable number of agents in pursuit of practical ends which are at least morally, practically, or otherwise normatively permitted (perhaps by relevantly “weighty” norms). Focusing just on the moral case for simplicity, if moral nihilism or some very revisionary moral realist view is correct and no or very few agents' practical projects ought to be thought of as morally permissible, then it could be argued that there is no (or not) enough instrumental value in the epistemic domain because there (either actually or simply in practice) is no relevant final value that epistemic states or doxastic attitudes conduce towards. Adherents of such skeptical and revisionist views may be best served by taking what follows to be hypothetical in nature, as an investigation of the non-skeptical, non-revisionist picture.

²² For more on the breadth of instrumental value in the epistemic domain, see chapter two.

²³ See Sylvan (2018).

to offer satisfying explanations, and the development of one's understanding. It is worth noting that the commonsense picture of this domain includes multiple positively evaluable epistemic states, though their relative fundamentality is open to argument.

While some philosophers may find the evaluations made from the epistemic point of view worthy of domain-independent acceptance in their final value theory, my suspicion is that others will either lack the sense that there is intrinsic epistemic value or, recognizing it, seek to explain this impression away, for example by reduction to another kind of value (perhaps practical value) or by offering an error theory. And it is worth noting that the adoption of the epistemic point of view is less a response to this ontological anxiety and more of a concession to it. My goal is not to overcome the skeptical or reductionist view directly in this piece but instead to offer an alternative picture with enough virtues to merit our attention.

Still, we should ask what the alternative reductionist or skeptical view about the intrinsic value of the epistemic domain might look like, for illustrative purposes. One simple and forceful version of this sort of attitude about the epistemic domain could be called an Instrumentalist attitude. An Instrumentalist would be willing to grant that there is epistemic value, but would hold that this value is wholly extrinsic--specifically, instrumental. It is easy to imagine an advocate of this picture: one need only imagine a deflationary Veritist.²⁴ Veritists accept that the fundamental epistemic aim is truth (often with some qualification, for example, truths about certain topics or truths relevant to certain agential aims). The Instrumentalist, as a deflationary Veritist, is in a position to explain their adherence to the view that the truth is of fundamental epistemic value by appeal to the instrumental, practical value of veridical epistemic states and doxastic attitudes. A

²⁴ For more discussion of Veritism see Berker (2013); cf. Goldman (2015).

careful Instrumentalist need not endorse the claim that veridical states are categorically valuable, and can instead emphasize the role that many or most veridical states have as sufficient for successfully getting about in the world--achieving our practical ends, and so on.

This seems to me a very sensible picture. But when this kind of instrumental value is taken to be an exhaustive account of epistemic value, it is difficult for me to avoid the impression that we are leaving out some important complications in order to preserve the simplicity of our view. There seems to me a clear and familiar sense in which certain agents can excel in a way that seems valuable without regard to downstream consequences, and that this sense is best captured by a non-deflationary view of epistemic value. This seems particularly salient if we imagine a pair of agents who end up equally well off in terms of the practical goods the Instrumentalist takes the epistemic goods to conduce towards, but who arrive at that position of parity through doxastic states that are different in ways that seem intuitively to *matter* epistemically--for example, if one of the two believes luckily.

These sorts of situations seem to me utterly mundane and familiar, the stuff of our everyday lives. We could try to capture the asymmetries in the different cases in terms of knowledge, or belief formation faculties or processes, mental capacities, skills, know-how, or via some other tool of philosophical analysis. Focusing on such cases inclines me towards the idea that enthusiasm for the hard-nosed Instrumentalist picture of epistemic value may be the result of some philosophers' lengthy focus on theoretically laden accounts of knowledge, and the (admittedly abundant) instrumental value of believing the truth.²⁵ I hope attending to familiar, mundane cases at least suggests that an alternative picture is

²⁵ For most of us, most of the time, given most aims. See section 2.2.

worth considering. I suspect that if such cases do not fill this motivational role, nothing else will.

This diagnosis, though I hope plausible, is of course not an argument that there is non-extrinsic epistemic value. Naturally, the subsequent work will appear better motivated to the extent that the intuition that there is such epistemic value seems forceful, and not easily explained away. Here is a schematic version of such an argument directed against an Instrumentalist. Returning to considerations introduced above, if we can specify two agents who are identically well-off in terms of the practical goods to which the relevant epistemic states conduce, but who nevertheless are not equally well-off in epistemic terms, then we can infer that there is some difference in value between the two agents which must be the result of epistemic value not captured by Instrumentalism.²⁶

I take it that such a case would be valuable for two reasons. First, it would not require that we identify some particular epistemic state or doxastic attitude as the state or attitude that is the epistemic difference-maker. This is an advantage since it allows us flexibility that the straightforward assertion of the value of e.g., knowledge, lacks. Of course, if we conclude that knowledge is the state that makes the epistemic difference, this will be compatible with how we have set things up. But it is to our advantage not to build the case against Instrumentalism in a way that presupposes that this one epistemic state is the bearer of distinctive epistemic value. The second reason to prefer the current approach to the stipulative one is because unlike the straightforward assertion of epistemic value, the presentation of an example, even a simple one, will leave us with a better sense of what it is to excel epistemically by providing a picture (albeit a limited one) of that kind of success.

²⁶ I am restricting the scope of this example argument to the non-instrumental rather than the entirety of the non-extrinsic in order to preserve simplicity. Considerations of more recherche varieties of extrinsic value, like indicator value, would require corresponding complications in the argument that, given the merely illustrative nature of this example argument, do not seem worth indulging in.

I do not mean to suggest that all epistemic success or value must be a close analogue of the kind given by the example I intend to discuss. The discussion focuses on an illustrative example, not a paradigmatic case to which all other cases of epistemic value are assumed to bear deep resemblance. This is important because it leaves the door open to views like those defended by certain Epistemic Virtue Theorists, who hold that there are distinct kinds of epistemic excellence (knowledge, wisdom, understanding, etc.).²⁷

One last proviso. In the sample case to follow, it is important to note that the epistemic states of the two agents themselves do not have to be identical, or even similar; we are challenging the idea that all value in the epistemic domain is extrinsic, specifically instrumental. (In this case, let us suppose for the sake of argument that this end is something like a rational agent's achieving some, perhaps constrained, practical goal or goals.) If the ends are equally well realized in two agents but there is nevertheless an evaluative difference between those two agents that seems best cashed out in epistemic terms, then we have good reason to take it that there is indeed non-instrumental epistemic value. I will now attempt to furnish such a case.

Imagine a near duplicate of a master carpenter. Suppose that this near duplicate behaves indistinguishably from the actual master carpenter during some (perhaps lengthy) episode of woodworking. The duplicate is different only because the duplicate's behavior is the result of a silicon chip implanted into their brain which forces them to, e.g., respond to such-and-such a question about woodworking or such-and-such a piece of knotty wood in just the same sort of way the actual master carpenter would. While the actual master carpenter's dispositions are grounded in a lifetime of practice and learning, the doppelganger's dispositions are the result of fantastical sci-fi interference.

²⁷ See e.g. Zagzebski (2001).

If the sci-fi chip-implanted-duplicate story seems offensively far-fetched, we can make our point sufficiently well with an amateur craftsman blessed with exceptional luck. Though this amateur does not know (or even truly believe) the same sorts of propositions about woodworking that the actual master carpenter does, and is in fact guided in their enterprises by false belief and guesswork, fate conspires to make the amateur's misguided attempts succeed in a way indistinguishable from the successes enjoyed by the actual master carpenter. When the actual master carpenter holds the chisel just so and, in light of their beliefs and dispositions based on experience with this kind of wood in similar conditions, hammers just hard enough to cut a mortise of the appropriate depth, the lucky amateur fumbles their way to an indistinguishable result, over and over again. When asked questions by apprentices the lucky amateur happens to guess--again, through sheer luck--the very responses the actual master carpenter would give (despite the lucky amateur's relevant false beliefs).²⁸ And so on.

It seems clear to me that neither luck nor the machinations of mad scientists make up the entirety of the epistemic gap between the doppelganger and amateur on the one hand and the genuine master carpenter on the other. Although the lucky performance of the amateur and the beliefs or belief-analogues of the doppelganger may suffice for all the practical goods that might follow from the master carpenter's epistemic states, there is a clear sense in which the actual master craftsman is better off than either the doppelganger or the amateur, and whether one prefers to explain that difference in terms of knowledge, know-how, or some other state, at least part of the difference is epistemic. I conclude that

²⁸ Those inclined towards certain dispositional accounts of belief may find it difficult to accept that an agent who so systematically acts against their belief in fact does hold the relevant belief. For relevant, although perhaps indirect engagement with this view, see section 2.3.

the actual master carpenter has more of the epistemic goods, and so I conclude that there are more epistemic goods to be had than merely instrumental ones.

Of course, the committed instrumentalist about epistemic value may deny this. They may suggest for example that this conclusion is a mistake due to our familiarity with the actual world, free of sneakily implanted carpentry chips and improbably lucky amateur woodworkers, and that if we were used to a world with such events, or if we simply took seriously the success rate of the silicon chip or the reliability of the amateur's luck, our seeming that the master carpenter is better off would evaporate. Perhaps it would. Like most arguments, this one admits of response. But in this case I do not find the response more compelling than the argument. And we have at least shown that there is a perspective ready to hand from which we can assess epistemic value non-instrumentally.

It might also be pointed out that our assumptions about the instrumental goods towards which epistemic states conduce were extremely sparse. It is easy to imagine this as the thin end of a dialectical wedge: the objector then goes on to suggest that if the breadth and diversity of the many personal and interpersonal outcomes of our epistemic states were adequately appreciated, the examples would fall apart. The bedrock response is that the Instrumentalist picture is objectionable because it is trying to assess epistemic states on the basis of their outcomes when it is clear that the connection between what goes on in the head and what happens subsequently outside of it are connected in many complex ways subject to many kinds of interference. The hard-nosed Instrumentalist accepts this and draws the conclusion that there are many ways to succeed epistemically: any way that results in practical success. The non-Instrumentalist will likely see this as a *reductio*, closer to an abandonment of the attempt to discover epistemic value than success in that endeavor.

It suffices for my purposes that we seem to have ready access to a point of view from which the actual master carpenter is epistemically better off than either pretender, and

that taking up that point of view does not seem woefully naive or metaphysically extravagant enough to merit dismissal out of hand.

1.3 Pritchard's Three Value Problems

Duncan Pritchard (2007) helpfully frames the existing literature about epistemic value by distinguishing several related issues that have been taken up by different writers. He calls these the *Primary*, *Secondary*, and *Tertiary Value Problems*. It's worth noting that unlike the Guiding Intuition, Pritchard's division of value problems is primarily meant as a guide to the issues that are taken up in the literature rather than primarily as a suggestion for how we ought to proceed. That said, I list them below.

Primary Value Problem: Why is knowledge is more valuable than mere true belief?

Secondary Value Problem: Why is knowledge more valuable than any [proper] subset of its parts?²⁹

Tertiary Value Problem: Why is knowledge more valuable than "...that which falls short of knowledge not merely as a matter of degree but of kind"?

Before discussing these briefly in turn it is worth noting the extent to which the existing literature, by generally focusing on versions of these issues, has focused on the relationship between knowledge and states such as true belief as opposed to the evaluative shape of the epistemic domain overall (perhaps better achieved by comparing knowledge

²⁹ "Proper" is my addition. Pritchard's use of "subset" invites a distracting objection to do with the nature of the subset relation that disappears once we add that he clearly has in mind proper subsets.

to states like know-how, wisdom, understanding, or features like epistemic virtues) and has taken for granted that knowledge *has* some special value when compared to e.g. mere true belief.

The *Primary Value Problem*, also sometimes referred to as the *Meno Problem* because of a tendency to identify the challenge in Plato's *Meno*³⁰ with the *Primary Value Problem*, is the question that many recent inquiries seem to have taken as their basic motivating question. If one sees themselves as investigating the *Primary Value Problem*, their goal is similar to vindicating the *Guiding Intuition*. While Pritchard's formulation directs our attention narrowly to a comparison with true belief, vindicating the *Guiding Intuition* would require us to say something about knowledge that would distinguish it from all other doxastic propositional attitudes. One distinction between the *Guiding Intuition* and the *Primary Value Problem* worth noting is that the 'thinner' one's preferred account of knowledge is--that is, the less one's preferred account of knowledge adds over and above true belief--the more perspicuous the *Primary Value Problem* becomes. Pritchard's narrow, contrastive focus on true belief seems to stack the deck against theories of knowledge that are in this sense thin.³¹

The *Secondary Value Problem* is usually brought up in the context of Gettier conditions on analyses of knowledge. Although the language here might seem to suggest a potentially dubious relationship of constitution, by "parts" Pritchard seems to mean only

³⁰ For more discussion of the challenge raised in the *Meno* see footnotes 45 and 53, but for now it is worth noting that the response Socrates offers ("episteme" is right opinion plus a "logismos aitiias") should strike many contemporary epistemologists, particularly those attracted by some form of Externalism, as very demanding if we accept that "episteme" ought to be translated as 'knowledge'. For discussion of this see Schwab (2015) and (2016). For more see Smith (1998) and Woodruff in *Epistemology: Companions to Ancient Thought 1*, Ed. Everson (1990).

³¹ This is worth noting especially given the defense by Goldman and others of knowledge ascriptions as genuine states of knowing where the truth conditions for those knowledge ascriptions are merely that the agent to whom knowledge is ascribed truly believe the proposition in question, see Goldman and Olsson on "weak knowledge", (2009). Compare Jäger (2014).

the various conditions knowledge must meet: that the proposition is believed, true, justified, satisfies some sort of anti-Gettier clause, and so on, according to one's preferred theory. Attention to this *Secondary* problem licenses questions of the following kind: if it is necessary for knowledge that it was reached safely, what role should the safety of belief formation play with regard to the value of the resulting piece of knowledge? Olsson and his interlocutors agree, as we shall see, that at least for Reliabilism, the Gettier condition ought not to be the condition that lends epistemic value. Riggs, on the other hand, seems to suggest that considering the role of the Gettier condition might be somewhere where we could make substantial progress.³² Pritchard's taxonomy allows us to cast this as a Secondary Value Problem dispute.

The *Tertiary Value Problem* represents a constraint on responses to the *Primary Value Problem* in that, if we take the *Tertiary Value Problem* to be well-motivated, we are thereby agreeing that only accounts of the value of knowledge that respond to the *Primary Value Problem* by showing knowledge to be more valuable as a matter of *kind* rather than of degree are acceptable. The relationship between the Tertiary Value Problem and the Guiding Intuition is intimate.

Theories that conceive of knowledge as true belief that rises to some justificatory threshold might seem to be in a difficult position with regard to the Tertiary Value Problem, since the move from mere justified true belief to knowledge on such accounts is just an accumulation of degrees of justification and not obviously a deep or explanatorily promising difference in kind.³³ One immediately available though perhaps unsatisfying response is to hold that knowledge is different in kind *just by* being true belief justified to

³² But see also Gardiner (2017).

³³ Someone who accepted a Williamsonian "knowledge-first" framework might be particularly well positioned to respond to the tertiary value problem, given the theoretical materials available to them: such views *do* make, indeed rely upon, there being an important difference in kind between knowledge and non-knowledge states.

the degree that meets the threshold for knowledge.³⁴ On this approach, meeting the justificatory threshold is not merely a change in degree of justification (or for example reliability of belief formation process type) but a change in degree of justification that *transforms* or *promotes* mere true belief into knowledge. This slippery thought might be extended as either a response to or rejection of the *Tertiary Value Problem*: the former by holding that this promotion lends or creates a distinctive kind of value, and the latter by holding that the availability of this response for accounts of knowledge that only issue in true belief plus changes in degree of justification undercuts any intuitive appeal the *Tertiary Value Problem* might have.

Importantly, Pritchard conceives of our task in responding to these challenges as admitting of some degree of what he calls “revisionism”.³⁵ Although he and others take each of these challenges to be well-motivated, it may turn out that an independently well-motivated theory of knowledge will fail to deliver a satisfying response to one of these challenges. In that situation we would have to weigh the other benefits of the theory against the motivations for accepting the value problem(s) to which it fails to respond. If we decided to accept a theory that failed to respond satisfyingly to one of the three value problems Pritchard poses, we would be being *Revisionary* in that we would be denying the importance of a challenge that carried intuitive weight.³⁶

³⁴ Proponents of this hypothetical line might be failing to take seriously the so-called Threshold Problem for knowledge (See e.g. Hannon 2017), but if one had a non-arbitrary reason for placing the threshold for knowledge at some level they would thereby be well-positioned to reply to the objection that the transition from mere true, justified belief to knowledge was a mere difference of degree rather than a difference in kind.

³⁵ Pritchard (2007).

³⁶ The Rawlsian notion of Reflective Equilibrium (RE) is helpful in grasping Pritchard’s notion of Revisionism. In RE we begin with some intuitions and other inputs, and consider whether conflicts emerge between them. We may, as a result of RE, abandon some intuitions that originally seemed forceful. Just so for Pritchard with respect to the various value problems.

1.4 A Dispute Between Frameworks

While Pritchard's taxonomy is helpful in coming to grips with the way many thinkers have oriented their inquiries into epistemic value, I prefer to focus on the Guiding Intuition. There are two reasons for this. One is that if we conceive of ourselves as investigating the three value problems directly and eventually conclude that we must be revisionary with regard to some of them, we will only have discovered that we set up faulty requirements on acceptable theories of knowledge. If we start with the *Guiding Intuition*, our task is instead to see to what extent we can vindicate an intuition with a lengthy philosophical pedigree and, more importantly, a great deal of intuitive attractiveness according to many contemporary epistemologists.³⁷ It is one thing to be revisionary about, for example, the Secondary Value Problem. It seems to me quite another to deny the Guiding Intuition. Revisionism about the former would just be learning that a question that appeared deep is actually shallow. (Perhaps composition is the wrong way to think about the relations between the features of an agent who knows some proposition, for example.) Revisionism about the latter would be a very surprising result; it would mean that our fundamental picture of the epistemic domain was badly formed in a surprising way.

A second reason to prefer framing our task in terms of the Guiding Intuition has to do with contemporary research into doxastic attitudes and epistemic states other than knowledge—for example, Understanding. Framing our inquiry into epistemic value as an inquiry into the value of knowledge, as much of the existing discussion already has, rules out from the outset the possibility of giving sustained attention to other doxastic attitudes. If it turns out that argument compels us to be revisionary about the Guiding Intuition, we would retain the option to move our attempt to locate epistemic value to some other doxastic state, rather than to conclude that there is no doxastic attitude of distinctive

³⁷ Grimm, Williamson, Kvanvig, Riggs, Goldman, etc.

epistemic value, just because the best going theory of knowledge fails to respond adequately to Pritchard's tripartite value problem. I think we naturally engage with knowledge because it seems the most promising way to vindicate a Weak Guiding Intuition:

Weak Guiding Intuition: There is some doxastic attitude that bears distinctive epistemic value.³⁸

I see the Guiding Intuition, as well as the literature surveyed by Pritchard about the value of knowledge, as an endorsement of the idea that we ought to begin our search for the doxastic attitude that bears distinctive epistemic value with knowledge.³⁹ Kvanvig, after concluding that there are serious difficulties in living up to his own account of the value of knowledge, suggests eventually that we might turn to understanding. I would characterize this maneuver as a retreat from consideration of knowledge as the best available way to satisfy the Weak Guiding Intuition. I take it that this characterization is informative, and I see the availability of offering this characterization as a reason to prefer conceiving of our task in terms of the Guiding Intuition.

The dialectical terrain is rocky here. There are surely some epistemologists who will either fail to see the appeal of the Guiding Intuition or, seeing some appeal, would not

³⁸ Although I am open to the possibility that a full inventory of properly epistemic value might well contain agential features or dispositions (such as open-mindedness) I'm going to restrict my attention to doxastic attitudes and epistemic states since I take it there is interesting work to be done here, and leave open questions about the kinds of agential features that are of interest to e.g. virtue epistemologists.

³⁹ This might strike some readers as exactly backwards. Such a reader might think we are only attempting to give a precise analysis of knowledge *because* it is antecedently clear that knowledge is valuable. Whether or not this somewhat Platonic idea captures the motivations of many contemporary epistemologists is unclear to me, but the more direct response to this approach is that it has proved fiendishly difficult to make the analyses of knowledge we have on offer live up to the Guiding Intuition. For my own part it seems clearer to me that there are valuable doxastic states than it is that those states are captured uniformly by any of these going analyses of knowledge.

be willing to abandon an otherwise well-supported theory of knowledge were it to fail to vindicate the Guiding Intuition. From such a point of view, any inquiry into epistemic value would likely seem otiose. There are likewise epistemologists so in the grip of the Guiding Intuition that they seem ready to dismiss any proposed theory of knowledge, no matter how well that theory fares with regard to other desiderata, if that theory fails to adequately capture their preferred view of epistemic value.⁴⁰ No doubt some such individuals will reject my proposing the Weak Guiding Intuition as a useful explanation of the central status of the Guiding Intuition.

I suspect it will be tempting for some to resist the replacement of the Guiding Intuition with the Weak Guiding Intuition by insisting that they are interested in (or disinterested in) the value of *knowledge*, not the value some other poorly understood epistemic state might bear. This is an important point. Such a complainant could point to Pritchard's Primary Value problem as evidence that other epistemologists are focused on the narrower task meant to be captured by the Guiding Intuition, and do not see themselves as engaging with the Weak Guiding Intuition in any direct way. I would like to agree that most discussion of epistemic value so far has focused on knowledge. I do not take either the Guiding Intuition or the Weak Guiding Intuition to be theses overtly considered and endorsed.

Here is the explanation I prefer for the evident preference in the literature for investigating the value of knowledge rather than a more general search for epistemic value. Recent epistemology was preoccupied for a lengthy period of time with an attempt to give a definition of knowledge in terms of necessary and sufficient conditions, almost to the

⁴⁰ Zagzebski seems to be an example of someone sympathetic to the latter viewpoint, when she proposes in *Recovering Understanding* (2001) that 20th century epistemology has gotten off-track in pursuing knowledge to the exclusion of wisdom and understanding.

exclusion of sustained investigation into other epistemic states. There are several important exceptions: Ramsey and others were interested in prediction, decision theory, and the scientific attitude of acceptance (as distinct from belief) in the early 20th century, and there has been some discussion of practical knowledge or “knowledge-how” on and off for several decades, but it is clear that the central topic within epistemology for a long period was the search for an analysis of knowledge.

When some contemporary epistemologists turned their attention to questions of value, it was within the context of this larger disciplinary dispute over the nature of knowledge. (After all, the Swamping problem was originally introduced as an objection to Reliabilism.) Although admittedly a piece of sociological speculation, it seems reasonable to infer that the preference for talk about the value of knowledge to epistemic value in general is the result of a preoccupation with knowledge within 20th century epistemology.⁴¹

Given recent work by Kvanvig and others to bring our attention to non-knowledge and yet distinctly epistemic phenomena, it seems to me methodologically preferable to acknowledge the Weak Guiding Intuition as the proper starting point, and, if one takes it that the proper object of investigation is knowledge, to offer arguments in favor of this position. If we end up accepting the view that the value of knowledge is to be the central question here, we ought to accept that view as the result of argument rather than assumption. It would be an embarrassing situation if our insistence on the primacy of knowledge was the result of a passing philosophical fashion, even a slowly passing one.⁴²

⁴¹ For further discussion of related ideas, see Zagzebski (2001) for an argument that modern epistemology as a whole represents a move away from philosophical interest in wisdom and understanding and towards knowledge and truth. Bonjour’s *The Myth of Knowledge* (2010) and Kaplan’s *It’s Not What You Know That Counts* (1985) are examples of more extreme arguments against the increasing complexity and artificiality of philosophical analyses of knowledge. See also Horvath (2016) and DePaul (2009).

⁴² The Primary Value Problem, or Meno Problem, is an interesting case study here. Although the Meno focuses on the value of “episteme”, the correct translation of that term is hotly debated among scholars of

Chapter 2: Truth, Belief, Knowledge, and Value

2.1 A Disclaimer About the Inquiry into “Mere True Belief”

There is a growing movement, pushed forward by Williamson and others, to build our epistemological theory not on the back of the notion of belief but instead with the notion of knowledge. On this “knowledge-first” view, knowledge is taken to be the “default factive attitude”⁴³ an agent takes towards a proposition. According to these theorists, knowledge is something in the light of which we ought to clarify other epistemic notions, as opposed to something best analyzed in terms of other, more fundamental notions (such as belief). On standard versions of this view, while it is true of agents who know that P that they also believe that P, their state of knowing that P is not constituted by a more fundamental⁴⁴ belief state which satisfies some other criteria, as on the traditional story.

Adopting the knowledge-first approach can make instances of mere true belief appear somewhat odd. Since the knowledge first theory fits naturally with (perhaps better: is frequently motivated by) the anti-skeptical thought that we generally get things right, and that generally those instances of getting things right count as knowledge, one might wonder what a mere true belief would be. Given a certain familiarity with or reliance upon an older paradigm, which we can call the belief-first framework (according to which epistemic theorizing uses belief in the role the knowledge-first framework reserves for knowledge) this thought can be difficult to appreciate. It seems natural to think that, given

ancient philosophy. While “knowledge” has been the traditional translation, several scholars are arguing for “understanding” as a better translation. See Schwab (2015). This is a minor dialectical point, but it would seem to take some of the wind out of the sails of the view that we ought to restrict our investigation to knowledge.

⁴³ For the canonical presentation of the position, see Williamson (2000).

⁴⁴ There is an important question about the contrast between explanatory and ontological nature of belief on the Williamsonian view. Although it is clear that the knowledge-first project is committed to at least an explanatory preference for knowledge, it seems natural to imagine this explanatory preference is explained by an ontological distinction. My discussion here presupposes this, but I don’t think anything vital hangs on interpreting Williamson this way at this point.

the belief-first view, even if the agent's knowing that P and believing that P are instantiated simultaneously, their knowledge *just is* a belief state that satisfies some other criteria. This makes the place of mere true belief easy to recognize: on the belief-first view, mere true beliefs are simply cases where the belief that P fails to meet whatever other criteria are required for a true belief to count as knowledge.

The knowledge-first theorist's task can seem more difficult by comparison. Since knowledge is the *default* attitude we take towards facts, it might seem that the knowledge-first theorist will have difficulties explaining how agents end up with mere true beliefs. Two points about this. First, this appearance likely stems from an unduly weighty reading of "default". Rather than understanding this piece of language as something like a disguised universal claim, I take it that if we understand it as a preponderance claim the difficulty vanishes, or at least is reduced to the difficulty of defending the idea that we often know and relatively rarely merely truly believe, a thought the knowledge-first theorist (and indeed any anti-skeptical epistemologist, whether knowledge-first or not) is likely to wish to vindicate eventually.

The second point is this: it's not that the knowledge-first theorist's theoretical *framework* requires her to say that mere true belief is any more or less common than the belief-first theorist's framework. But what it does illuminate is that the knowledge-first theorist's picture fits more naturally with a conception of mere true belief on which such states are degenerate pieces of knowledge. The knowledge first theory seems to suggest that, in general, agents end up with mere true beliefs as a result of *failed attempts to know*. In other words, in order to end up merely truly believing, something has (at least usually) gone wrong: either the evidence is worse than one took it to be, one is treating their evidence irresponsibly, some strange set of unlucky (or lucky!) circumstances obtains, or some other epistemic problem or defect is present. For my own part, I think the knowledge-

first characterization of mere true belief is more attractive than the characterization typical of the belief-first view. When I try to imagine cases of mere true belief, instances of agents who truly believe that P but are such that it seems inappropriate to attribute knowledge of P to them, they often have an unusual or defective character just as the knowledge-first approach seems to predict they will.⁴⁵ Many of the most natural cases of this type are charlatans or epistemic ne'er-do-wells of one stripe or another: lucky guessers, improbable tarot card readers, fallacy-inclined gamblers, and so on.

We should take seriously the assumptions built into the framework we use to discuss mere true belief when inquiring into the value of that state, especially if we take seriously the idea that the value of a state or object is not simply an additive result of the values of the parts or components of that object. If a mere true belief is, as a belief-first theorist might say, not best thought of as flawed in some important sense but instead merely as a believing that has not achieved the lofty status of knowledge, this may well have an impact on our thinking about the value of mere true belief when compared to an alternate picture on which, as the knowledge first framework seems to suggest, mere true belief is (at least generally) the result of something's having gone wrong somewhere along the line in an attempt to know something. Most importantly for our purposes, on a view like the knowledge-first view, it becomes more natural to regard mere true belief as generally *less valuable* (in the many cases where it is the product of a flawed process) than we would likely be inclined to take it to be on the traditional belief-first view.⁴⁶

⁴⁵ With the possible exception of cases of putative mere true-belief introduced as part of a theoretically laden discussion of certain contextualist views, motivated by investigations of “high-stakes” epistemic scenarios often taken to motivate Contextualist views of knowledge.

⁴⁶ Spelling out the sense in which such lucky or misguided mere true beliefs will lack value compared to other possibly benign mere true beliefs will be difficult, for reasons that will become apparent in section 2.2. Despite the ability of the lucky guesser's guess to, e.g., guide them successfully to Larissa, it seems clear that “from the epistemic point of view” something has gone very badly wrong in such cases.

For this reason, it seems possible to me that the ascendancy of the knowledge-first approach could signal a sea change in the way we think about the value of mere true belief. This is especially important given the dialectical role of analyses of the value of that state. Much existing work comes to conclusions about the value of knowledge by contrasting it with the value of true belief. Recall, for example, Pritchard's approach to epistemic value—specifically his *Primary Value Problem*. The Primary Value Problem explicitly directs us to compare the value of knowledge to the value of mere true belief. We need to be careful in assessing our background assumptions about true belief before making this comparison.

If the belief-first theoretical framework is apt to mislead us about the value of mere true belief, we have some reason to be skeptical of theories of the value of knowledge that rely on the belief-first framework and are argued even partially on contrastive grounds. It seems possible to me that much older work will need to be reconsidered in the light of this and other insights afforded by the knowledge-first framework. Having flagged this possibility, in subsequent sections, especially in chapter three, I wish to assume the orthodox, belief-first picture in order to more naturally engage Goldman and Olsson.

2.2 Knowledge and Instrumental Value

As I've stated it, the *Guiding Intuition* is very weak, and deliberately so. It demands only that knowledge be shown to have value that stands out when compared to the value borne by other doxastic attitudes. This is because I take it that many doxastic attitudes will bear one kind of value at least some of the time: instrumental value. Plausible accounts of knowledge should yield the result that agents who know that P are well-positioned to achieve related practical ends. If a proposition, T, states the location of the nearest train station, then in most cases (given moderately cooperative circumstances) agents who know

T and aim to arrive at that train station are well-positioned to do so. If they succeed by making use of their knowledge, it will have been of instrumental value to them.

It is tempting to conclude that this instrumental value is typical of knowledge, and then to ask whether this kind of value could be the value noted in the *Guiding Intuition*.⁴⁷ There are two responses worth offering here. The first is typically taken to motivate the *Primary Value Problem*: this instrumental value, insofar as it seems to generally flow from the truth of the thing known, seems to be achievable without satisfying the further conditions on true belief⁴⁸ that would promote it to knowledge. As in the *Meno*, it seems that knowledge's instrumental value does not distinguish it from mere true belief.

To clarify, if Jane is trying to get to the nearest train station and merely truly believes that T, but succeeds by acting in accord with this true belief, it seems this true belief has been just as instrumentally valuable as the corresponding piece of knowledge would have been. Jane would be like Socrates' man on the road to Larissa.⁴⁹ And this would make it the case that we had failed to vindicate the *Guiding Intuition* by showing that knowledge has instrumental value, because this instrumental value would not be distinctive of knowledge, since it is shared by mere true belief. The *Guiding Intuition* is not merely that knowledge is valuable, but that it at least tends to have some kind of *distinctive* value that makes it stand out when compared to other doxastic states.

⁴⁷ Again, living up to the *Guiding Intuition* does not require showing that knowledge is categorically valuable.

⁴⁸ Williamsonians may prefer to avoid talk of "adding" conditions to true belief in order to achieve knowledge, given that on their view it is knowledge itself that is the "default" doxastic factive attitude, but on their view there is still an important difference between knowing and merely believing, so I take it the point is merely linguistic.

⁴⁹ Kaplan (1985) and Zagzebski (1996) mention that the Socratic project preserved in Plato, as related to "episteme", might not be best thought of by analogy to knowledge but instead to something like understanding. What to make of this suggestion given the unclarity involved in the latter notion isn't obvious, but given that this exegetical claim about the *Meno* is usually serving in a motivational role for a question of obvious independent interest, I will leave the exegetical issue aside for now. See Schwab (2015) for more discussion.

The second point regarding the instrumental value of knowledge has to do with the instrumental value, not of mere true belief, but of some false beliefs. As noted above, knowledge *generally* has instrumental value, but not in all cases. Sometimes knowledge is a hindrance. There are many pairs of agents and practical aims such that, for some proposition Q, an agent is more likely to succeed given that they falsely believe \sim Q rather than truly believing (or knowing) Q.⁵⁰ An example may help make the structure of a class of these cases clear.

Suppose Tom wishes to make the U.S. Olympic swimming team. Suppose also that given Tom's times in his chosen events, it is unlikely that he will make the team. Call this true proposition, that *Tom is unlikely to make the team*, Q. A final stipulation: if Tom believes (or knows) Q, he will lose motivation to train, making it even less likely that he will go on to represent the U.S. in the Olympics. Plausibly, if Tom were to believe \sim Q, this would not only avoid the deleterious effects on Tom's efforts to realize his goal and to join the Olympic team that would accrue were he to believe (or know) Q, but could actually help motivate Tom to continue training. Tom could falsely think to himself: *I am well-positioned to succeed, so long as I continue training*. Given this picture of the motivational upshot of belief in \sim Q, it seems that Tom's false belief might well be of great instrumental value, given his goal.

The general outline of these cases of instrumentally valuable false belief are clear enough: an agent is situated with regard to some practical aim such that for some relevant proposition they are more likely to succeed given a false belief than a true one. Some might object that mere likelihood of success does not suffice for these beliefs to produce real instrumental value. The details of instrumental value generation are not perfectly clear but

⁵⁰ And, perhaps importantly, we are loath to attribute some sort of defect to the agents involved. The features of their psychological makeup that enable false beliefs to play the right kind of role seem utterly common and benign.

this objection seems fair enough. Happily we do not need to take a strong position on the nature of instrumental value to respond satisfyingly. Against this objection we need only slightly modify the case such that the agent does succeed in their aim given their false belief, and that had they believed the relevant true proposition, they would have failed. Though perhaps less frequent, cases with this structure do not seem objectionably unusual or uncommon. So long as there are any cases with this character, we may conclude that instrumental value is not unique to knowledge nor to mere true belief, but instead that even false beliefs can be of instrumental value.

One way to think about what has gone awry in these cases is to think about one schematic (broadly Humean) view about practical reasoning, according to which practical reasoning requires an affective or conative component, as well as a doxastic component. In effect, proponents of such views take intentional action to require both a representation or picture of the world and a related goal or aim--a map and a compass, if you like. Though knowledge is one doxastic attitude on which one could act,⁵¹ one could also act and succeed on the basis of a mere belief, or perhaps even some other doxastic attitude, like a suspicion or prediction. Though others have discussed cases where instrumental value results from true and even false beliefs as well as knowledge, I haven't yet seen the point made that this is to be *expected* given that at least on a broadly Humean view of practical reasoning, it is possible to act on the basis of any of a multitude of different doxastic inputs--neither knowledge nor mere true belief is uniquely suited to the role of doxastic input into the practical syllogism.⁵² Section 2.3 will briefly expand on this issue.

⁵¹ Perhaps even the norm of action, as discussed in Williamson (2000) and the subsequent literature.

⁵² Although the Humean theories of motivation and action are often challenged, it is typical of these challenges that it is the conative or affective state that is regarded as superfluous. Most challengers seem ready to grant that there is a doxastic component to intentional action, so my point here should not fall prey to this familiar criticism of the broadly Humean story of motivation and action, see e.g. Dancy (2000).

Whether or not this appeal to practical reasoning to explain the potential instrumental value of different doxastic attitudes is illuminating, not only is instrumental value not *distinctive* of knowledge (as shown by Jane), it is not even guaranteed by knowledge (as shown by Tom). More generally, while it seems that most knowledge may be instrumentally valuable, this does not seem to be a promising approach to the *Guiding Intuition*. There we were looking for some valuable feature of knowledge that was distinctive among competitor doxastic propositional attitudes. Instrumental value has turned out to be poorly suited to that role. So where should we turn?

One possibility is to look elsewhere within the realm of extrinsic value—specifically, by considering the possibility that knowledge is distinctively valuable because it bears *indicator value*. On this approach knowledge is not valuable in itself, but neither is it merely instrumentally valuable. Instead, knowledge is distinctively valuable because unlike other doxastic states, it indicates the presence of something finally or intrinsically valuable. Just what this further bearer of intrinsic value is and why knowledge indicates it while other states do not remains to be seen, but this is the approach favored by Olsson. Discussion of this approach will occupy chapter three.

2.3 Acting Without Knowledge (Or Belief)

What kinds of doxastic states can feed into our practical reasoning? The formulation of some extant theories of practical reason in terms of belief (or knowledge) gives rise to a natural question about the possibility of acting on the basis of non-belief doxastic states.⁵³ Specifically, is there reason to restrict the doxastic states involved in successful or appropriate practical reasoning to belief or knowledge? Is it plausible to

⁵³ Williamson (2000) introduces discussion of the knowledge norm of action. See also Schechter (2017) for an evaluation of the related Williamsonian knowledge norm of assertion.

exclude credal states? Formal epistemologists or those who are ecumenical in their epistemological theorizing, welcoming both credences and beliefs, might well find hesitation over the possibility of acting on the basis of a credal state a puzzling kind of timidity, or perhaps a relic from an earlier era of theorizing that predated contemporary formal epistemology. They might be right. While credences are likely the most perspicuous case of a plausible doxastic state that adherents of views such as the Knowledge Norm of Action omit, the debate over beliefs and credences seems to me too fraught to be of great help when it comes to answering our question about acting without belief.

Instead of focusing on credences, we can ask another question: can an agent act on the basis of a mere acceptance?⁵⁴ Suppose a physicist is asked to make a prediction about the position of a thrown object. This physicist *knows* that a simple model which ignores, for example, drag, will be to some extent inaccurate for that very reason (suppose there is no compensatory mechanism for the force the theory glosses over). Nevertheless, despite her zero credence in the (perfect) accuracy of the model, she may, on the basis of the flawed model, predict the object's later position with enough precision for some particular practical purpose (e.g., winning a bet). The prediction to which the model guides her, though false, suffices for practical success.

She may go on, if questioned, to say that she acted on the basis of a model she knows to be inaccurate, based on a theory she knows to presuppose several false oversimplifications, but she nevertheless *accepts* the theory for the purpose at hand. These sorts of cases seem to me to be the most straightforward problem cases for the view that agents cannot act on the basis of mere acceptances. The physicist does not merely decry believing that the object will be precisely where the theory predicts it will be, but is in the

⁵⁴ As noted above, Patrick (1990) lays out the relationship between Van Fraassen's notion of "acceptance" and belief.

stronger position of *knowing* that the object will not be (exactly) there, since she knows the theory on which the prediction is based is false (and false in a way that predictably precludes an accidentally true verdict in this case). So her nevertheless acting in accord with the theory (or, more awkwardly but more cautiously: acting indistinguishably from how one who did believe the theory would act) seems to show that it is possible for agents to act without the doxastic component of their action being a ‘full’ belief.

It’s possible that the judicious application of some sophisticated epistemological machinery may avoid this sort of counterexample. One might reply that she reasons on the basis of her *belief* that the model is adequate for her purposes, or her belief that ‘the object will be very close to where the theory predicts it to be’. Still, while she may well have either or both of these beliefs, it is unclear that they are operative in her practical reasoning. Further recourse might be made to a dispositional theory of belief on which, at a rough pass, the agent’s disposition to act as if the theory is true *just is* their belief in the theory. Such an approach promises to collapse this family of alleged counterexamples.

Either way, the upshot is clear. Tokens of any type of doxastic attitude or epistemic state which can feed into practical reasoning can be of instrumental value. It seems to me reasonable, for the kinds of considerations just raised, to interpret the class of doxastic attitudes and epistemic states that can play this role broadly. Whether one is convinced of this or not, it is clear that the broader one takes the class of (permissibly) practically actionable epistemic states to be, the less distinctive their instrumental value will be. If knowledge’s value is to be at least partially distinctive, as the *Guiding Intuition* indicates, this strongly suggests, although does not prove, that we ought to look beyond instrumental, practical value in our attempt to satisfy the *Weak Guiding Intuition* by appealing to knowledge.

2.4 Mere True Belief, Veritism, and Epistemic “Junk”

At this point we’ve seen reason to conclude that true belief often but not always has instrumental value insofar as it helps agents accomplish their plans and satisfy their desires. We may still ask if it is plausible to hold that mere true belief has other non-instrumental value. Schematically, we are turning from trying to explain the value of true belief in terms of its instrumental value or nature as a kind of success to examining the possibility of holding mere true belief to be an intrinsic good (at least from the “epistemic point of view”). Is it intrinsically good to (merely) believe the truth? Given a certain understanding of *Veritism*, the view that the fundamental good in epistemic inquiry is the truth,⁵⁵ it might seem plausible to respond affirmatively.

One intuitive problem with this kind of view was put forward in a relatively well-known way by Harman and later discussed influentially by Sosa, and Friedman, among others.⁵⁶ Harman and others offer examples of agents who deliberately collect epistemic “junk”: allegedly valueless epistemic states. The older discussions focus on junky knowledge, typically, but it is not hard to imagine junky true belief.⁵⁷ Imagine someone sitting alone on a long sandy beach on a pleasant day. They begin counting the grains of sand on the beach, building up their catalogue of knowledge. *There are X grains of sand on this beach*, this person silently affirms, then *there are X+1 grains of sand on this beach*, and so on. This case is similar in some ways to an idle disjunctivizer, who, having come to believe some truth, begins mentally affirming disjunctions of the true thing they believe with far-flung possibilities.⁵⁸

⁵⁵ For more discussion, see Berker (2013) and Goldman (2015).

⁵⁶ See footnote 19.

⁵⁷ That is, if we set aside considerations to do with the nature of “mere true belief” discussed in section 2.1.

⁵⁸ Thanks to Daniel Eaton for this example.

One way to describe the strangeness of these kinds of behaviors relies on the (frequently challenged) distinction between implicit, or tacit, belief and explicit, or occurrent belief. Both agents are doing something odd in making what are at least normally recondit implicit beliefs occurrent. Surely for at least small values of X, the person sitting on the lonely beach already had sufficient (perceptual) evidence to warrant believing that the beach had that many grains of sand, and indeed if they had been asked, for example, whether or not the beach had five grains of sand on it, we should expect them to respond affirmatively and so attribute to them *at the outset* a dispositional belief with that content. Plausibly their perceptual evidence suffices to justify their occurrent belief to the threshold necessary for it to count as knowledge, barring any unusual additions to or stipulations about the case. So I think there are some difficulties involved in both the beach case and the idle disjunctivizer case, for parallel reasons. It's just not clear that they are in fact *gaining knowledge* when they direct their thoughts in the ways described.

At this point I will turn from trying to frame our inquiry into epistemic value to looking at one particular response to the *Guiding Intuition* on behalf of *Process Reliabilism* offered in a recent paper by Goldman and Olsson. Although other responses and supplementary suggestions have been offered,⁵⁹ Olsson's preferred response, the *Conditional Probability Solution*, is the most promising approach on offer for the Process Reliabilist, and has received the most critical attention.

⁵⁹ For example, Goldman's own preferred response, "Value Autonomization", from the same 2009 paper, but see also the exchanges between Olsson and Jäger.

Chapter 3: The Swamping Problem and the Conditional Probability Solution

3.1 The Swamping Problem

Consider a toy version of reliabilism on which an agent S knows some proposition P if and only if the following three conditions are met:⁶⁰

1. P is true;
2. S believes P; and
3. S's belief that P was produced through a reliable process.

One suite of related objections to reliabilist accounts of knowledge, called the Swamping Problem, has to do with the relationship between the truth of the thing known and the reliability of the process by which S comes to know it.⁶¹ The swamper's general strategy is to start either by assuming or arguing that knowledge has some important kind of value. They then try to show that the value of mere true belief is, or very nearly is, the same as the value of reliabilist 'knowledge'. (Hence the name: the value of the veridicality of the belief state *swamps* any value reliability might add.) Swampers conclude that reliabilist 'knowledge' isn't actually knowledge at all, since reliabilist 'knowledge' is not valuable enough to be *genuine* knowledge.

Here is how Swinburne puts the issue. By "the reliabilist requirement", he means the central feature of reliabilism, captured above in our sketch of the view by condition (3):

⁶⁰ This is a slightly modified version of the model view Goldman and Olsson (2009) consider, omitting the anti-Gettier clause per their subsequent discussion.

⁶¹ Jones (1997), Swinburne (1999), Zagzebski (1996; 2000; 2003), Kvanvig (2003), and others have offered related concerns as well, see Goldman and Olsson (2009) for further examples.

Now clearly it is a good thing that our beliefs satisfy the reliabilist requirement, for the fact that they do means that...they will probably be true. But, if a given belief of mine is true, I cannot see that it is any more worth having for satisfying the reliabilist requirement.⁶²

Zagzebski's discussion makes the nub of the issue even more clear. She holds that, "[T]he reliability of the source of a belief cannot explain the difference in value between knowledge and true belief..."⁶³ To see why, she offers a helpful analogy with espresso machines. Suppose we work in an office with two espresso machines, one which reliably produces tasty espresso and which only does so unreliably. If you ask me to get you a tasty espresso and I do so, there is a sense in which it does not seem to matter whether it came from the reliable or the unreliable machine. As Zagzebski sums up, "...If the espresso tastes good, it makes no difference if it comes from an unreliable machine...If the belief is true, it makes no difference if it comes from an unreliable belief-producing source".⁶⁴

The proponent of the swamping problem suggests that succeeding with regard to tastiness *swamps*, or overwhelms, the value of the reliability of the source of the espresso, if indeed the reliability lends any value at all. And so the story goes with regards to truth and reliability, according to Zagzebski, Swinburne, and others.

Helpfully, Goldman and Olsson are very clear about their understanding of the challenge they intend to respond to, and characterize it this way:

What Zagzebski is saying is that the value of a good espresso is not raised by the fact that it was produced by a reliable espresso machine *if taste is all that matters*; and, likewise, the value of a true belief is not raised by the fact that it was produced through a reliable process *if truth is all that matters*. [Italics original.]⁶⁵

⁶² Swinburne (1999), page 58.

⁶³ Zagzebski (2003), page 13.

⁶⁴ Ibid.

⁶⁵ Goldman and Olsson (2009).

Here Goldman and Olsson emphasize a somewhat unusual feature of the standard interpretation of the swamping problem, on which reliability is taken not to contribute any value at all. (Their explanation of this is that Zagzebski's analogy reveals that she is assuming that truth is all that matters in inquiry.) Although this is not essential to offering a swamping-type challenge, framing the issue this way simplifies matters immensely. Before embracing this simplification it is worth noting that several extant versions of the swamping problem are precisified in ways that build in the idea Goldman and Olsson draw attention to, that reliability, in the end, adds *no* value to a true belief.

Specifically, Goldman and Olsson write that, "Once truth is in place, its value appears to swamp the value of reliability, thus making the combination of truth and reliability no more valuable than truth itself".⁶⁶ There doesn't appear to be a temporal reading of "once" here that makes sense. If we read "once" above as meaning something like 'given that', Goldman and Olsson's view seems to be that whatever value reliability may add does not survive contact with the truth of the belief. They, and others, seem to think of swamping in terms of deletion or cancellation. It is possible to hold that the value of a complex state is not always equivalent to the sum of the values of the constituent parts of that state, but principles like these are neither mentioned nor invoked.⁶⁷

There's at least some temptation to interpret claims of the form 'Consideration A *swamps* consideration B' in the same way we would interpret claims of the form 'Consideration A *overwhelms* consideration B', with the implication that B had some weight to it. When we explain our decisions by saying things like, "my promise to help my friend move outweighed my desire to watch the game", we have set up the issue in a way that seems to acknowledge a countervailing force. The precisifications of the Swamping

⁶⁶ Ibid.

⁶⁷ Although see Allen (2003) for a discussion of Moore on organic unity.

Problem we will look at it *lose* this part of the challenge to reliabilism, presumably in order to preserve simplicity. And while simplicity is valuable, it is worth noting that given the swamper's strategy as outlined above, this simplification is for them purely a convenience. The swamper could equally well argue that the value added by reliability, while non-zero, does not suffice to grant reliabilist 'knowledge' the advantage in terms of value over true belief we ought to expect. (In Swinburne's terms, this approach would acknowledge that our beliefs are *only slightly* more worth having for satisfying the reliabilist's requirement.)

The simple version of the swamping problem that Goldman and Olsson consider, which holds that reliability adds no value to a true belief, seems to have sacrificed some of the intuitive appeal of the swamper's objection. It simply isn't necessary to get the objection running that one hold that reliability adds *no* value to true belief. It suffices for the objector's purpose if the value added by a reliable process of belief formation is so trivial that the value added by truth *swamps* the value added by reliability, where "swamps" is interpreted to mean something like 'dramatically outweighs'. Reliabilist 'knowledge', according to such an objector, would not be knowledge at all, faring only marginally better than mere true belief in terms of value, and so, by appeal to something like the *Guiding Intuition* or perhaps the intuition behind Pritchard's *Tertiary Value Problem*, inapt as an account of *knowledge*.

Of course the attack would be cleaner if the swamper could show that reliability is truly valueless, or if one was willing to assume that only truth matters in inquiry, as Goldman and Olsson take Zagzebski to do. The swamper would not then be occupying the middle ground, holding on the one hand that reliability is valuable in at least some epistemically relevant sense, while on the other withholding from states that bear this value (and which meet the other relevant criteria) the title of knowledge. The *prima facie* awkwardness involved in holding this view does not seem to me to amount to a serious

objection against it though, and it preserves both the intuitive force of the swamper's objection (that there is not enough distance between reliable true belief and mere true belief) as well as the weighty body of intuitive evidence suggesting that reliability of belief formation often simply suffices as justification for knowledge (these would be the intuitions suggesting that in non-demon world, non-Norman style cases, simple reliabilism seems like a satisfyingly explanatory account of knowledge). Whether these considerations lead one to prefer the more complex version of the swamping problem or not, it is worth pointing out that there is another way to run the objection.

Swinburne's remark in particular seems to me helpful in understanding how some swampers see the key issue. He seems to suggest a comparison between two cases along roughly these lines. We shall hold fixed that my belief that P is true in both, only adjusting the etiology of that belief. In one case it is the result of a highly reliable belief formation process, and in the other, it is the result of an unreliable one. Swinburne asks what the sense is in which I am better off in the former case than the latter. One straightforward way for me to be better or worse off would be for me to have some information about my epistemic states. If I were to know that I knew that P, or came to the belief that P reliably in the former case, or if I were to learn that I came to believe P flukily in the latter case, it seems clear that I would be better positioned to put my epistemic state to practical use in the former case than the latter. In the former case I might become more confident that P, or my belief that P might become more resilient, more resistant to certain kinds of defeat. In the latter I might even abandon my true belief that P, or my belief that P might simply become modally weaker, easier to dislodge by rebuttal or the application of sophistical pressure.

These all seem like relatively straightforward ways in which I might be better or worse off in one or the other of Swinburne's pair of cases. But Swinburne tells us that he "cannot see" how his belief is better for being reliable, which strongly suggests that he is

focusing narrowly on the belief itself, and excluding these sorts of meta-level epistemic states. Like other epistemic states, etiological reliability isn't 'luminous': we are not infallible detectors of the reliability of the process that led to our beliefs. Perhaps if reliability were luminous, the swamping problem would not be compelling. But without some kind of direct access to the facts about the reliability of the process that led to the formation of the belief in question, we are left wondering whether two individuals with doxastic states that are indistinguishable from an internal perspective differ in terms of value. Swinburne here seems at least tempted by a negative response. Although the two cases differ dramatically in the kinds of evaluations we should make from a third-personal point of view (after all, according to the reliabilist, one is a case of knowledge while the other a mere true belief), for the subjects involved there is no detectable difference--there is simply the belief that P.⁶⁸ So it seems strange to hold that the one state is of greater value than the other while focusing narrowly on the beliefs themselves and not meta-level epistemic states or the distant consequences of the two states, as Swinburne seems to do.

Another related way to press a swamping style worry has to do with some familiar considerations about the nature of reliability itself. Reliability comes in degrees, but the justification it confers is either sufficient for knowledge or not. If we assume that we can give relatively fine-grained analyses of the reliability of the processes that led to the adoption of our beliefs (especially plausible if we construe reliability as a long-run historical notion although not implausible if we take reliability to be best understood counterfactually) then we should be able to compare beliefs formed by the most reliable *unreliable* process and the least reliable *reliable* process. Though we can imagine a reliability theorist settling on an arbitrary stipulative threshold in order to have a fully

⁶⁸ As before, I am setting aside Williamsonian questions about the nature of the relationship between belief and knowledge.

determinate theory--suppose 80 percent, 76.5 percent, or 51 percent reliable, this detail is irrelevant--it is difficult to accept that such a stipulative threshold marks a bright-line distinction in the *value* of the belief produced by a process that barely meets the reliability threshold when compared to one that only very narrowly misses it.

3.2 The Conditional Probability Solution

In responding to the swamping problem as their opponents put it to them, Goldman and Olsson take their task to be showing that knowledge has some *extra* value over and above that of true belief:⁶⁹ they label the challenge to which they hope to respond the EVOK, or Extra Value Of Knowledge. This is noteworthy in that they only aspire to show that knowledge has *more* value than true belief, not an importantly different kind of value as some (e.g., Kvanvig, Pritchard's *Tertiary Value Problem*) have hoped for, nor a great deal more value (as it would take to satisfy the more sophisticated version of the Swamping Problem just discussed).

Their response to the swamping problem is an attempt to show that knowledge satisfying clause (3) above, what Swinburne called "the reliabilist requirement", gives us a state that *exceeds* the value of true belief just in terms of the value of truth. Although the extra value Goldman and Olsson locate for reliabilist knowledge is subtly different from the kind of value possessed by mere true beliefs, it too eventually reduces to a matter of apprehending truths. This means their response doesn't require them to assert that, (and so does not require us to consider whether) reliability per se has any special epistemic value of its own, thus avoiding the main thrust of the Swamping problem altogether.

⁶⁹ For a helpful taxonomy of other challenges to the value of knowledge, see Pritchard (2007), discussed at length in sections 1.3 and 1.4.

They call their response to the swamping problem the conditional probability solution, and it is admirably simple. Given the simplified reliabilist view stated above in (1)-(3), an agent who knows that P came to that knowledge by using a reliable belief formation process. Given that some “empirical regularities”⁷⁰ are in place, that agent’s coming to know P per reliabilism means that they deployed a belief forming method that is likely, if deployed again, to yield future true beliefs in relevantly similar circumstances.⁷¹ As Goldman and Olsson put it, “...the probability of having more true belief...in the future is greater conditional on S’s *knowing* that P than conditional on S’s *merely truly believing* that P” (28).

This feature of reliabilist knowledge, according to Goldman and Olsson, gives that state Indicator Value. As mentioned above, Indicator Value is a kind of extrinsic value borne by objects or states when those states indicate the presence of some other object or state which is itself valuable. In this case, token beliefs have indicator value because they indicate the presence of a reliable belief forming process. Given the way we’ve set up this toy version of Reliabilism, it will be true that any case of putative reliabilist knowledge will indicate a reliable belief forming mechanism. But is this mechanism valuable?

Goldman and Olsson’s opponents have granted that true belief has value, as is plausible. Given this, it would be at least awkward to hold that true belief is valuable but

⁷⁰ For example, that the belief forming method is not unusually unique, can in principle be used again in the future, is such that an agent is likely to deploy it again, and so on. See Goldman and Olsson (2009), and section 3.3 for further discussion.

⁷¹ There are serious open questions for even sophisticated reliabilist theories of knowledge to do with belief-formation process individuation and evaluation, usually called the generality problem for reliabilism. Conee and Feldman’s “The Generality Problem for Reliabilism” (1998) is the classic statement of the issue. See Goldman’s “Epistemic Folkways and Scientific Epistemology” (1993) for the makings of an agent-reliabilist approach to a potential solution. Although there are several current proposals being worked on as responses to the generality problem, none is yet widely accepted. Still, the fact that there are open problems for reliabilism does not preclude its being valuable in its role here to suggest the extent of one plausible externalist response to questions of epistemic value.

that the prospect of getting more true beliefs in the future is of no value.⁷² The swampers Goldman and Olsson engage also deny that reliabilist justification a la clause (3) above grants any value in itself. The conditional probability solution, as Goldman and Olsson are quick to point out, is “silent” on this point. Their CPS requires only this slight extension of the idea that true beliefs are valuable in order to get off the ground. Still, we can ask whether or not this maneuver is a satisfying response to Swinburne’s complaint. Does this recondite extrinsic feature of a particular reliably-formed belief make *the belief itself* “more worth having”, or have we instead found that it is better to have the belief?

3.3 Evaluating the Success of the Conditional Probability Solution

The conditional probability solution seems to me to succeed as far as it goes: specifically, it seems to succeed as a response to the EVOK problem. This is not to say that the conditional probability solution is deeply satisfying. It carries with it several important limitations. Although Goldman and Olsson may have succeeded in responding to the task they set themselves—to show that a simplified reliabilist view *can* show that knowledge (although only in a generic sense) has more value than true belief, given some plausible assumptions—their doing so falls far short of what other philosophers have suggested we look for in an account of knowledge’s value.

One cause for concern is that the conditional probability solution only vindicates what Goldman and Olsson call a “generic” value for knowledge. By this, it is meant that

⁷² This is not to assert a false piece of value theory, viz., that adding more of a valuable item always results in more value. Although that principle is arguably false in general, it is unclear why we would think it would fail to apply in the relevant case. One might have a Harman-style concern about accruing valueless “junk” true beliefs, but the kinds of cases that makes Harman’s thought most plausible have to do with an agent, for example, idly disjunctivising as discussed in section 2.4. This is at least not typical behavior, and as we shall soon see Goldman and Olsson do not defend what Kvanvig (2003) has called a “categorical” account of the value of knowledge but instead only a weaker “generic” account. For Harmanian type issues to arise at this point one would need to think that most belief acquisition is (or could be in relevant cases) junky. This is implausible.

most instances of knowledge carry the value that they think EVOK tells us to look for. On their view, some reliabilist knowledge is such that the agent's coming to know some proposition does not increase the chances that they will end up with future true beliefs, either because the prompt to which they responded in coming to believe was particularly unusual, or the method they employed to form the belief in question was unlikely to be used again in the future. In these cases, reliabilist knowledge, even by Goldman and Olsson's own lights, is no more valuable than mere true belief. This means that the CPS is not going to appeal to theorists (Kvanvig, etc.) who seek a resolution to the question of epistemic value on which knowledge is always, as opposed to merely usually, valuable.

Goldman and Olsson are conveniently explicit about their goals. Their EVOK problem calls for only extra value, not value of a particular different *kind*. Some epistemologists, attracted to what Pritchard calls the Tertiary Value Problem, have offered that the deep challenge we face is to show why knowledge is not just more valuable, but valuable in a different way than other doxastic attitudes or epistemic states, such as, true belief, are valuable. Kvanvig discusses the idea that knowledge is *categorically* valuable; that is, important in a way that is distinctive when compared to other doxastic attitudes. While reliabilist knowledge, per the conditional probability solution, is valuable for reasons that differ from the reasons for which true beliefs are valuable, these different reasons only have to do with the likely acquisition of future true beliefs. There is only one item of deep, *fundamental* value reflected in the proposal Goldman and Olsson offer, and that is believing the truth. So while it is true that the conditional probability solution shows that knowledge has added value by having all the value of believing the truth as well as value that derives from an extra, less direct connection to truth via the reliable belief formation

faculty,⁷³ it seems unlikely this will satisfy philosophers who sought out a sense in which knowledge's value is distinctive or categorical.⁷⁴

Further, because it is offered in response to the EVOK problem, the conditional probability solution is only responsive to one kind of swamping challenge to reliabilism. As discussed, on the version of the Swamping challenge pushed by Zagzebski et al., *all* the value of knowledge is explainable in terms of the value of truth. Whatever other conditions knowledge carries, including justification or an anti-Gettier clause, can add no value. But as discussed above, a second understanding of the Swamping challenge is possible, and to my mind more natural. On this version, what we mean when we object to Reliabilism on the grounds that the value of truth *swamps* the value of knowledge is that the gap between the value of a mere true belief and the value of reliabilist knowledge is *too small*—not that there is no difference between their values. Truth's value swamps the value of e.g. justification or anti-Gettier conditions by being overwhelmingly greater, not because those features contribute nothing once combined with truth.

On this second version of the swamping problem, the swamper does not deny that reliabilist 'knowledge' is more valuable than mere true belief, but instead holds that while there is a difference, it is altogether too small to vindicate the intuition that *knowledge* holds a place of prestige in terms of epistemic value. While Goldman and Olsson are keen to read Zagzebski and other swampers as primarily interested in truth and truth-conduciveness, this might be a mistake. If the swampers are really looking for something other than reliability which *does* make a belief more worth having, they may well be looking for just

⁷³ Kvanvig (2003) makes a relevant point about value and mereology that seems to have been overlooked by Goldman and Olsson (and others). Kvanvig rightly insists that it does not suffice to show merely that knowledge is a composite state with multiple valuable parts. We must further offer independent reason to think that these values will complement, rather than interfere with or detract from, each other. Sylvan's discussion (2018) takes up this issue.

⁷⁴ It straightforwardly abandons the categorical value claim, though we have already seen some reason to doubt the categorical value of knowledge due to Harmanian junk-type concerns in section 2.4.

the sort of broadly internalist features Kvanvig notes offer a promising approach to the value of knowledge, such as the ability to defend, explain, or offer reasons in defense of one's belief.

Relatedly, the extent to which the indicator value that makes up the heart of the conditional probability solution is a satisfying response to Swinburne's question about the "belief itself" is unclear. On Goldman and Olsson's analysis, the feature that turns out to bear the extra value we sought is an extrinsic feature of the belief. Arguably the belief *itself* has not been shown to exceed the value of true belief; instead, what we have learned is that some beliefs are connected to a new perhaps unexpected source of epistemic value--belief formation faculties of a certain kind. And it seems clear that it is these underlying features that are directly responsible for the indicator value reliabilist knowledge has been argued to (generally) have. So it seems to me possible that if one was gripped by Swinburne's concern as exposed above, and was willing to narrowly focusing on the "belief itself" and not the downstream consequences or extrinsic properties of the belief, one might argue that we have not really shown knowledge to have any extra value at all. Rather, we've made an adjacent discovery about the value of some belief-forming faculties.⁷⁵ Instead of a vindication of knowledge, we've instead shifted our attention away from doxastic attitudes altogether towards agential capacities of a certain sort. While this does seem to suffice as a response to the EVOK problem as Goldman and Olsson set it up, this can seem to say more about their framing of the problem than their resolution of it. This approach would require some delicate footwork, since it seems to force us to deny that indicator value adds value, but it is available.

⁷⁵ For a parallel argument about the value of modal features of knowledge (safety and sensitivity), see Gardiner (2017).

Whether one takes the strident response just discussed to the CPS or not, it seems unlikely that the CPS will satisfy swampers despite solving the EVOK problem, because the swampers are after more than just another unusual kind of extrinsic value. Kvanvig, for example, wants a view on which knowledge has categorical value. Zagzebski is interested in a view that ties knowledge more closely to virtues like wisdom. Due to their mistaken reading of their interlocutors as primarily focused on truth, Goldman and Olsson's work simply does not engage with the deeper considerations that seem to be motivating their critics. To invoke Pritchard's terminology, Goldman and Olsson's narrow focus on responding to the Primary Value Problem by showing that knowledge has a bit of extra value isn't dialectically effective because their interlocutors seem to accept the Tertiary Value Problem. Before turning to my positive proposal, there is one more point I'd like to make about the CPS and its limitations.

3.4 Framing the Generality Problem

The Generality Problem is worth attending to in this context not in order to draw attention to serious, although familiar, difficulties for Process Reliabilist theories of knowledge, but in order to show that the Generality Problem bears on the discussion of the value problem, and in particular the CPS, in ways that so far are undiscussed. I think the criticism outlined in section 3.3 suffices to show that the CPS is unsatisfying as a response to the value problem, but here I will introduce the issues I see surrounding the Generality problem as an auxiliary line of attack. The upshot is that the CPS is at best incomplete: its current presentation runs roughshod over a crucial distinction, and thereby presupposes a particularly strong response to the Generality Problem.

Per Reliabilism, to find out whether some agent S is justified in believing that P, we look to the process that led the formation of the belief that P. If that process is reliable

(and some other conditions hold, for example, an anti-Gettier condition) then S is justified in believing P. Although this proposal is admirably simple and commands significant intuitive force, careful philosophical attention has revealed that this story glosses over an important distinction. Beliefs are not delivered *ex nihilo*; they are formed as a result of processes at least some of which, plausibly, are or supervene on physical processes in the brain of the believer.⁷⁶⁷⁷ This process (or group of processes) produce the belief. Following Conee and Feldman, we can call it (or them) the Token Belief Formation Process. Although it may seem ungainly to refer to tokenings of processes, this is not worth muddying the waters over. We can take the name to indicate (i) their tight connection with the token belief into which we are inquiring, and (ii) the fact that we will later need to speak of these Token Belief Formation Processes as instances of more general types of processes. Although the terminology might be thought to be unfortunate, all that is required is that when we inquire into a token belief, we accept that some particular process (of brain activity, mental activity, or both) led to its adoption. The particular series of events that led to the adoption of such a token belief is all that the term ‘Token Belief Formation Process’ is meant to indicate.

There is an important sense in which Token Belief Formation Processes are not a good fit for the Reliabilist machinery. Reliabilism tells us that when we are wondering whether a belief is justified, we ought to proceed by assessing the reliability of the process that led to the production of that belief. We have just seen that token beliefs will be the

⁷⁶ It is tempting to add that these physical processes lead to a mental event or cause (or are) a transition of thought that leads to the adoption of some belief, but some prefer to think of beliefs as bare dispositions and would regard this addition as misled. Whether one prefers a story including more mental or physical elements does not affect the core of the Generality Problem, only the way it is set up.

⁷⁷ Klein’s discussion (1999), inspired by Sosa’s account of the “kinds” of knowledge we should attribute to thermostats, which do not have brains, should I think be taken metaphorically, and so not construed as a serious objection to the supposition here that the processes that lead to the formation/adoption of a belief are, partially are, or supervene on, brain states, neural activity, and so on.

causal result of Token Belief Formation Processes (TBFPs). These TBFPs, though, will not give rise to interesting reliability facts. We are asking about some *particular* series of synaptic events and/or mental transitions. As “tokens”, or, less clumsily, individual processes, we can say about them only that they led to the adoption of a veridical or a non-veridical belief. If forced to respond to a question about their reliability, the most we could say of the TBFPs, as they are necessarily singly-instantiatable, would be that those that led to veridical beliefs are perfectly reliable, having led to true beliefs, and of the TBFPs that led to non-veridical beliefs that they are perfectly unreliable, having led to false beliefs.⁷⁸ A view that took the processes that the Reliabilist theory assesses to be TBFPs would surely be a monstrous misconstrual of Reliabilism, and would result in devastating extensional errors, assuming only that one countenances justified false beliefs or (even less contentiously) unjustified true beliefs. To make sense of the insight motivating Reliabilism, we need to embed the TBFP within a broader category or type, something multiply instantiated, so that this new, broader thing could give rise to the kinds of interesting reliability facts the view plausibly instructs us to seek out.

Put differently, TBFPs will look like, for example, some neural activity in an agent followed by the adoption of a belief. This isn't the sort of thing the Reliabilist wanted to make justificatory facts a function of in the first place. Luckily, a natural response suggests itself. We can embed Token Belief Formation Processes into categories called Belief

⁷⁸ Alston's (1995) distinction between counterfactual and long-run or historical reliability does not help us here. In assessing a TBFP we are evaluating this or that particular series of events. The long-run or historical reliability of these processes will only ever include their single deliverance. The counterfactual question, what would have happened if an agent were to form a belief on the basis of *this very* TBFP, only seems to result in the agent's forming a belief with a content that matches the content of the belief they in fact formed. I do not take this to be an alarming determinism but instead a straightforward result of the fine-grainedness of the TBFP. To expect the counterfactual reliability of a TBFP to vary is just to accept that *this very* chain of synaptic activity and so on could have resulted in a different belief. This seems to me implausible.

Formation Process Types (BFPTs).⁷⁹ Unlike TBFPs, These BFPTs will include just the sorts of things Reliabilism enjoins us to look to when searching for something to underwrite the reliability facts.⁸⁰ Perception, visual perception, visual perception in normal lighting, assessing one's evidence, and so on will all be BFPTs.⁸¹ These are all examples of BFPTs into which Token Belief Formation Processes can be sorted. And because BFPTs can be multiply instantiated (modally or long-run/historically, whichever one prefers), we can fruitfully ask how reliable a Belief Formation Process Type is. Unlike TBFPs which can only return reliability results of 0 or 1, BFPTs can return results from 0 to 1 inclusive, just as the Reliabilist would've hoped.⁸²

Hopefully everything to this point seems like a friendly, if somewhat unwieldy, clarification. The simple statement of Process Reliabilism above required us to assess

⁷⁹ Again, following Conee and Feldman.

⁸⁰ There is some controversy over whether to construe the objects of reliability analysis modally or non-modally, by taking processes (Process Types, in the language of Conee and Feldman) to be individuated counterfactually or long-run/historically. Alston offers a compelling argument that the long-run/historical approach opens the Reliability theory up to bootstrapping concerns, but it is not clear whether the modal approach encounters problems of its own. Whichever option one prefers, the Generality Problem will still have teeth.

⁸¹ When Goldman introduces the Reliability theory the types of processes he takes the theory to assess are very revealing; he takes it that justification will be a matter of whether or not a token belief is the result of for example, "good reasoning". It seems to me that this feature of his presentation is responsible for much of the intuitive appeal of the view. It can sound almost definitional to say that a belief is justified if it is the result of good reasoning. I am somewhat concerned as to whether or not it is good practice to admit types of this intuitive kind without making their extension more clear by making their meanings more precise, but it seems to me the proponent of the Generality Problem can admit them without affecting the argument. More would need to be said to relate the Agent/Virtue Reliabilist-like view Goldman proposes in his "Folkways and Scientific Epistemology" to these issues. Such a view can forestall the Generality Problem by limiting the number of BFPTs. A proponent of Generality-like concerns might have to preclude the Agent/Virtue Reliabilist from listing their Agential or Virtuous BFPTs on the grounds that either (i) the short list of BFPTs the Agent/Virtue Reliabilist offers is arbitrary, or (ii) that those BFPTs cannot be individuated in a principled way from extensionally indistinguishable non-virtuous/agentially virtuous BFPTs. (e.g., believing P because one used good reasoning sounds virtuous, but believing P because is instantiating one of many disjuncts in a lengthy disjunction of all the instances captured by "good reasoning" doesn't sound virtuous. Our inability to disambiguate between these in a principled way just is the Generality Problem, on my view, as explained below.)

⁸² This should not lead us to consider abandoning TBFPs, though: while TBFPs cannot give rise to interesting reliability facts, BFPTs cannot be causally efficacious. It is the tokens, not the types, that must cause beliefs, but the types, and not the tokens, that must be assessed for reliability.

processes for reliability and to use these to generate justificatory verdicts with regards to beliefs. This statement overlooked an important point to do with the difference between a belief's proximate causal process, which cannot itself interestingly be assessed for reliability, and a larger group of relevantly related causal processes, which can. Fair enough: the Reliabilist need only look to the latter. But everything needed to run the Generality Problem has now been established. We need only note that each Token Belief Formation Processes will instantiate many--indefinitely many--Belief Formation Process Types, and that these BFPTs will vary wildly in their reliability. So for one and the same belief, Reliabilism is left in the awkward position of seeming ready to return many contrary verdicts. Further amendment is clearly required, and the general direction of that amendment seems clear. We need to add to Reliabilism a mechanism that selects which BFPT to assess for reliability. The Generality Problem is the task of giving an account of this selector.

3.5 The Conditional Probability Solution and the Generality Problem

Goldman and Olsson posit four “empirical regularities” that need to be in force in order for their Conditional Probability Solution (CPS) to show that knowledge has indicator value. Their claim with regard to these regularities is not that they are always met, but that they are satisfied at least some large portion of the time, a large enough portion to make the claim that “knowledge is valuable,” understood as a generic rather than universal claim, true (or whatever passes for truth in the case of apt generic claims). These four regularities are called *non-uniqueness*, *cross-temporal access*, *learning*, and *generality*. Although Goldman and Olsson are not explicit about the distinction between Token Belief Formation Processes (TBFPs) and Belief Formation Process Types (BFPTs), I will use these notions to precisify their brief remarks, discussing each of these regularities in turn

and showing why Goldman and Olsson see them as limiting the range of cases in which they take the CPS to assign indicator value to knowledge.

Non-uniqueness has to do with “problems” agents face, where a problem seems to be something like a situation that prompts the formation of a belief. The idea here is that the CPS relies on the possibility of instantiating a BFPT multiple times. If the agent fails ever to form a belief, Q, as a result of a TBFP that instantiates the same BFPT as some temporally earlier belief, P, then the agent, S, in knowing P, is not in a state that has indicator value. This follows because on the Conditional Probability Solution, S’s knowing that P is valuable because this knowledge is indicative of S having recourse to some BFPT that S can employ again. If S forms the belief that R in response to a unique problem, even if S comes to know that R, this state will not have any indicator value in the sense CPS assigns, because S will never have another opportunity to employ a TBFP that instantiates the BFPT that lead to⁸³ their belief that R, by, the uniqueness of the problem that led to the formation of R. So non-uniqueness limits the scope of the CPS by telling us that knowledge formed in response to unique problems has no indicator value.

Cross-Temporal Access tells us that S’s knowledge that R is only valuable in the sense picked out by the CPS when the TBFP that S employed in coming to believe R is an instantiation of a BFPT that can in principle be “accessed” or instantiated again in the future. If, for example, we are considering whether or not S’s knowledge that R is valuable and we assume that this knowledge is the result of a Token Belief Formation Process that is an instantiation of a Belief Formation Process Type that is somewhat like Φ_2 above, except that this process type is in principle unable to be instantiated in the future again, then the CPS will not hold that S’s knowledge that R has indicator value. Where

⁸³ As the above section attempted to show, making this claim seems, in a strict sense, disallowed. We lack a principled way to point to *the* BFPT that a given TBFP instantiates in the way that matters for Reliabilist justification.

Uniqueness ruled out cases of agents confronting “problems” they would not face again as cases where CPS would assign indicator value, Cross-Temporal Access rules out cases where an agent is confronted with a non-unique problem, but solves it by recourse to a BFPT that they cannot in principle “access” again--by using a ‘unique’ problem solving method, so to speak. Without being able to use the BFPT⁸⁴ that led to an agent’s knowing that P again, CPS cannot assign their knowledge that P any indicator value.

What further restrictions should we place on the CPS? So far, Goldman and Olsson have told us that the CPS assigns indicator value to Reliabilist knowledge when an agent solves a non-unique problem by a method that they can in principle access at a later time. The next “regularity” or “assumption” is Learning. They give the learning regularity in this form, where S is an agent and M is a “method”⁸⁵ used to solve a problem: “By the learning assumption, S is likely to make use of M again [when S is confronted with a problem of the same type that they used M to resolve earlier].” The Learning constraint, then, extends the CPS by requiring that for a piece of knowledge to have indicator value it must not only arise as a result of an instance of a non-unique problem solved in a cross-temporally accessible way, but that the knowledge be the result of a kind of cross-temporally accessible method that an agent is likely to use again just when they are confronted with other problems of a similar type.⁸⁶

⁸⁴ As above.

⁸⁵ “Methods” as Goldman as Olsson use them are ambiguous between TBFPs and BFPTs in Conee and Feldman’s language. Put differently, Goldman and Olsson are adding the CPS to just the sort of Reliabilist view that Conee and Feldman complain elides the distinction between process tokens and process types.

⁸⁶ There is a confusing exchange in a pair of papers between Olsson and Jager on just this point. Jager seems to accuse Olsson of having conceded in correspondence that Olsson took the Learning constraint to be pseudo-internalist: in order to satisfy the learning constraint, the agent *needs to learn*, in some sense, the method used to arrive at the belief. This is a significant departure from orthodox Reliabilist externalism, and an unnecessary complication to the CPS. All Goldman and Olsson need is that the agent is *in fact* likely to redeploy their problem-solving method, or perhaps that they in fact will redeploy the method, not that they need access to their belief formation processes or to bear any attitude towards them. So “learning” is probably a misleading name for this criterion.

The final regularity is Generality, which is introduced as follows, where M, again, is a method used to solve a problem: “By generality, M is likely to be reliable for solving that similar future problem as well”. So it seems generality is just the assertion that methods for solving different problems of the same type are likely to be reliable for all problems of that type given that they are reliable for solving any problem of that type. Whether Generality is best understood as a regularity in its own right or as a complication of Non-uniqueness isn’t perfectly clear. One might think that in Generality we are just getting a clarification of the somewhat opaque term of art “problem”.⁸⁷

At this point we have the entire set of restrictions on the scope of the CPS that Goldman and Olsson lay out. Knowledge has indicator value just where S knows that P on the basis of a method M that responds to a non-unique problem, M is in principle accessible for S in the future, S is likely to re-deploy M when confronted with future problems of the same type, and M is likely to be reliable when used in connection with those future problems. I see two primary issues here. First, as stated, the empirical regularities Goldman and Olsson discuss presuppose a perspicuous solution to the Generality Problem. Second, their presentation of these regularities belies a lack of attention to the distinction between TBFPs and BFPTs.

⁸⁷ It’s also worth noting the strength of the particular claim Olsson and Goldman make. They do not claim that a method M that is reliable with regard to solving a particular problem will be a reliable method for solving future problems of the same type, but instead that M is *likely* to be reliable for future problems of that type. By the other constraints outlined above, the CPS would assign indicator value only where an agent confronts a non-unique problem by employing a method they could in principle use again and are likely to redeploy. But Generality, as stated, guarantees us only that this method is likely to be reliable. So it could be unreliable. But this would mean that in these cases, where M is in fact unreliable, the CPS would assign indicator value to a belief-forming method that will not lead to future knowledge. This seems contrary to the spirit of the proposal and the apparent intention of the other three constraints on the CPS Goldman and Olsson lay out. It’s possible that this stipulation, that a method would likely be reliable, rather than the assurance that it would be reliable, is motivated by concerns about the individuation of problem types, perhaps for reasons similar to the type individuation problems surrounding the Generality Problem.

Tackling the first issue first, why is it that these regularities presuppose not just a solution to the Generality Problem, but a strong solution? In addition to their role in constraining the valuable kinds of knowledge, these four criteria might be thought of as, or related to, desiderata on a solution to the Generality Problem. These are just the sorts of features of a BFPT that might make it appealing, or that might highlight a BFPT as the sort of type that features in our pretheoretical picture of thinking and reasoning. But stating these constraints does not offer us progress towards a solution. It is well and good to say that we might want a solution to the Generality Problem to highlight the kinds of processes that would satisfy the empirical regularities Goldman and Olsson stipulate--to be accessible cross-temporally and general to some reasonable standard. But this is not the difficulty we face in confronting the Generality Problem. It is not that we lack plausible BFPT candidates, but instead that *selecting* among the various BFPTs that a given TBFP instantiates requires solving a very difficult problem to do with the naturalness (and, very likely, the relevance) of process types. So to say that we should look to, for example, cross-temporally accessible process types, while a reasonable thought, is not a suggestion that promises progress with regard to the the crux of the Generality Problem. We already knew we wanted an answer with that feature. But there will be indefinitely many BFPTs with that feature that any given TBFP will instantiate.

The second difficulty is that much of what is said by Goldman and Olsson elides the distinction between TBFPs and BFPTs. “Methods” are ambiguous between TBFPs and BFPTs in just the way that the careful work of Conee and Feldman, and others, should’ve warned us against. To understand Goldman and Olsson’s constraints, particularly Cross-Temporal Access and Learning, we must read methods as BFPTs, but their work seems to rely on the idea that for any one belief there is only one or only one operative BFPT at play, as opposed to the unique TBFP that led to the adoption of that belief. And this is just the

assumption that the Generality Problems reveals to be misleading. Not only is there always more than one BFPT into which a given TBFP can be sorted, *none* of the BFPTs are operative in the sense of being directly causally responsible for the adoption of a token belief. Only TBFPs stand in this direct causal relationship to beliefs. So the infelicity of arguing that we can “use” a belief forming method, where ‘method’ is read as a BFPT, reflects a lack of attention to the crucial distinction between Token Belief Formation Processes and Belief Formation Process Types.

I have already discussed concerns stemming in a more direct fashion from the shape of the CPS. Here I have tried to push a complementary line of criticism by bringing the Generality Problem into discussion of the CPS. The CPS only assigns indicator value to knowledge that is the result of “methods” that meet certain criteria. Unfortunately, without a solution to the Generality Problem, stating these criteria seems difficult, and finding unique BFPTs for a given belief such that we could find Reliabilist knowledge to assess by the lights of the CPS seems impossible. Again, my aim has not merely been to rehearse well-known problems for the Reliabilist research program, but to point out that one of these issues in particular throws up a significant second roadblock for proponents of a response like the CPS to Reliabilism’s Value Problem.

Chapter 4: Understanding

4.1 Kinds of Understanding

Having completed the negative component of the project by suggesting some difficulties involved in the Reliabilist's attempt to convincingly explain epistemic value, I'd like to turn now to my positive proposal. While Goldman and Olsson show that there is a way for the Reliabilist to show that their analysis of knowledge has more value than mere true belief, it isn't clear that they have satisfyingly responded to the Primary Value Problem. If one takes the Tertiary Value Problem seriously, it is likely that their approach will seem wanting. This leaves us in a difficult position. We could abandon Reliabilism as a theory of knowledge on value theoretic grounds, as Zagzebski seems keen to. We could also try to explain away the epistemic value intuition, as Goldman seems to suggest we should.⁸⁸ Neither approach is fully satisfactory. As an alternate proposal, I'd like to suggest we finally consider a move proposed in chapter one: retreating from the Guiding Intuition to the Weak Guiding Intuition. Specifically, I'd like to consider the possibility that Understanding is the state best suited to capture what Kvanvig and others hoped to recover from knowledge.

"Understanding" is in wide use. It has moral, linguistic, and epistemic uses, at least. Here we are interested in the epistemic uses of the term. It'd be appealing if we could offer a persuasive, unified account of understanding in its epistemic guise, but most attempts so far seem to focus on either propositional or objectual understanding. I am interested in Understanding because I think it is a familiar, important feature of our epistemic lives, and one that affords an opportunity to sketch a relation between an agent and a content that is sophisticated in the way that the most promising 'thin' theories of knowledge are not. If

⁸⁸ See footnotes 18 and 60.

we can find a different propositional attitude that recovers much of what the swampers want to say about knowledge, we may be able to hold up these epistemic achievements without upsetting work in the theory of knowledge.

There are familiar non-epistemic uses of understanding, as well as epistemic uses that differ in terms of their superficial structure. While some claims take the form “S understands *that P*,” others take a different form, where understanding takes a (non-propositional) noun phrase as its complement, for example, calculus. I’ll call the first *propositional understanding* and the second *objectual understanding*. The straightforwardly propositional uses are much better understood. There are promising analyses of these on offer in terms of knowledge.⁸⁹ An analysis in terms of knowledge is very appealing, as an opportunity to leverage our existing epistemic machinery. The received view of propositional understanding does just that, while borrowing from other sources as well. On one version of this view, to understand that P is to know why P. On another, understanding P requires knowing some R, such that R explains P. These views are built to capture the idea that understanding is knowing *why*. What is it to understand that it is wrong to burn cats for fun? One tempting view is that this understanding just is knowledge why recreational cat-burning is wrong. Despite the lively debate over the best version of this view of propositional understanding, it is unclear how to apply this theory to the objectual uses.

One cause for concern for the propositional cases of epistemic understanding is that their analysis in terms of knowledge seems to open the door to cases of junk (propositional) understanding. Whatever is objectionable about the person on the lonely beach counting up grains of sand seems to be objectionable about someone who accumulates

⁸⁹ See e.g. Grimm (2014).

understanding by affirming to themselves that, for example, *this* sock is in the dresser because I put it there after washing it, and *this* sock is in the dresser because I put it there after washing it, and so on.

While proponents of such views of propositional understanding are keen to remind us that on their view understanding is a special kind of knowledge only insofar as it has a special subject matter, and not in virtue of there being additional criteria which need to be satisfied, this seems to leave their proposals insufficiently distinct from knowledge to serve as the bearer of distinctive epistemic value in lieu of that state.

There is less of a developing consensus about objectual understanding. Most discussions seem to start with the idea that the epistemic state of interest here relates an agent to a set of propositions, individuated by something like relevance. Theories diverge over whether or not an agent can understand some domain if their knowledge of (the relevant part of) the set of propositions involved is only partial, or if some of their relevant beliefs are false. It's helpful to think about a more concrete kind of case. How much of which epistemic goods (knowledge, justification, true belief, etc.) does a chef need to accrue, and with regard to which propositions, before it becomes appropriate to attribute understanding of, e.g., baking, to that chef? Is one false belief about baking sufficient to preclude the chef's understanding?

To sum up our survey thus far, I'll offer a small chart that takes advantage of two fundamental distinctions in order to break up the logical space. With these in view, it is clear which kinds of theory are well-suited to which linguistic data. It also dramatizes the reasons why a simple, unified analysis of epistemic understanding would be surprising. We can distinguish singularly propositional from non-singularly propositional epistemic states by appeal to the nature of the content those states relate agents to. We can also

distinguish factive and non-factive attitudes by appeal to the propositions that can stand in the relevant relations.

Unsurprisingly, the received view of propositional understanding sits in the top left quadrant of our chart, as this view reduces understanding to knowledge, and so analyzes propositional understanding in terms that are factive. Minority views according to which it is possible to understand a false proposition lie in the lower-left quadrant. The right column contrasts two approaches to objectual understanding: on one view, all the propositions in the set agents are related to must be true. On the other this constraint is not in place.

	(Singularly) Propositional	Non-Singularly Propositional
(Directly) Factive	Received view of propositional understanding	Objectual understanding of sets of facts only
Non-Directly Factive	Propositional understanding of falsehoods possible	“Robust Objectualism” ⁹⁰

Table 1: Theoretical approaches to epistemic understanding.

4.2 On Objectual Understanding

In this section I’ll sketch the view of objectual understanding that I prefer, and note some difficulties with the project. At this stage, the topic is ripe for inquiry, and although the general approach to the objectual understanding seems promising and well-motivated, a completed analysis in terms of necessary and sufficient conditions is not forthcoming, though some possible necessary conditions will be discussed. I do think we have enough of a proposal together to see the general shape of the theory of objectual understanding and appreciate why it would serve well in a partnered role as a state responsive to the Weak

⁹⁰ The term here is due to Khalifa (2017).

Guiding Intuition with a theory of knowledge like Reliabilism, but the details of a theory of objectual understanding are unsettled, for now.

On my view, objectual understanding is a doxastic propositional attitude, but it is not a singular propositional attitude. Understanding in this sense is a relation between an agent and a topic or content construed as a set of believed propositions of cardinality at least greater than one. I do not endorse a truth requirement for membership in the set, or an ordering requirement on which some special subset of propositions in the group must be true, justified to some threshold, or meet other requirements. Some sufficiently high percentage of the members of the set must be true, and there are good reasons to think that there is a content-based restriction on which subjects or topics can be understood. On my view agents who understand a subject generally ought to be able to defend their relevant beliefs, manifest an associated skill or capacity, or explain the object of their understanding, although we can furnish counterexamples to each of these constraints. Finally, I believe that such a view of objectual understanding hews closely to the ordinary concept of epistemic understanding in its objectual use, and is well-suited to satisfy the Weak Guiding Intuition as the doxastic attitude that bears distinctive epistemic value. In the remainder of this section I will try to explain and defend each of these points in turn.

First, the non-singularly propositional nature of objectual understanding. Recall that we are interested in the state attributed by utterances like “Jane understands calculus.” Although it is possible to defend a view on which there is just one conjunctive or otherwise complex proposition that Jane understands, her relationship to which captures her understanding of calculus, this view seems like a mere curiosity. If two agents each understood calculus but differed doxastically would we then be forced to adopt a polysemous account of “calculus”? Besides, such an utterance can be perfectly apt even of agents who make occasional errors, coming to false beliefs about calculus through the use

of their existing mathematical belief formation methods. It is far simpler to imagine Jane's understanding as a matter of her bearing a group of doxastic attitudes to a group of propositions rather than taking her to bear one attitude to a massive proposition that encapsulates all her calculus thoughts while also bearing contrary attitudes to other contents, as in the case of false belief due to error.

Should we require that all the propositions in the set be true? Scientific cases seem to suggest that we shouldn't. Surely Newton, in developing his physics, improved his understanding of motion. This is close to a paradigm case of objectual understanding. Nevertheless, given what we now know about the Newtonian mechanics, any too zealous adherence to that theory at least skirts endorsing falsehoods. For a less radical example, consider the master woodworker discussed earlier, or the chef mentioned just above. Surely one false but relevant belief does not suffice to overturn one's understanding of a subject matter. For preface-paradox type reasons, any factivity constraint would threaten to eliminate objectual understanding altogether, or reduce it in extension to the extraordinarily doxastically cautious. And this would do a great deal of violence to the commonsense concept of understanding we are out to capture.

This next point is delicate. Given that it seems we must allow *some* false beliefs to co-exist with objectual understanding of a subject, we need some means to limit the objects of understanding. If we don't, we risk being in the awkward position of trying to hold up this particular epistemic state as being specially valuable but without a way to rule out tarot card reading, astrology, and other obviously bankrupt enterprises. Two strategies suggest themselves. First, one could attempt to give a characterization (or a laundry-list) of the features of a subject or topic that preclude it's being understood in the laudatory sense we're currently interested in. A second approach would be to try to locate as a necessary

condition some central feature, perhaps some perspicuous propositions within the set of propositions understood, with a character that ruled out the inappropriate topics.

I'm skeptical that we will be able to furnish a non-ad hoc list of forbidden topics or that the appropriate objects of understanding are structured in a way that would make finding similar central propositions plausible. Perhaps one needs to know something about the chain rule in order to understand calculus. But objects of understanding can be singled out at varying degrees of fineness of grain. Just as one might understand baking, one might understand breadmaking, or sourdough particularly, or the history of sourdough in a certain region of France in the eighteenth century. As we shift from one scale to another, it seems likely that the propositions that are central to an agent's understanding of related topics will shift. This suggests that locating an agent-independent central core, even for one topical area, will be difficult, and so seems to suggest that any injunction to use a core content to determine whether or not a subject or topic is understandable is misguided.

Whether we understand this as support for the idea that we ought to adopt a laundry-list approach to the extent of understanding or not, there does seem to be one feature that rules a topic out: sufficiently deep or widespread inconsistency. Certainly the Newtonian mechanics can be understood, false though they may be, but it is difficult to imagine someone understanding that vaccines cause autism. Any body of belief that contained a sufficiently large proportion of falsehoods might be impossible to understand for just that reason. But it seems important that genuine understanding requires more than just a group of related true beliefs, or even bits of knowledge. They need to be related to each other in the right way. Understanding generally seems to be connected with abilities like the ability to explain, defend, or demonstrate. While such abilities are far too psychologically demanding to feature in plausible accounts of knowledge, they seem to be the very markers of understanding.

None of these capacities seems necessary in order to understand a subject, though. While the general connection between understanding and ability seems relatively robust, it is easy to picture someone who understands a concept but, perhaps due to an illness or some other dysfunction, is unable to put their doxastic attitude to practical purposes. Though one might understand perfectly well why this piece of gnarled timber will not serve for a given project, they could be incapable of expressing this understanding. So it would seem to be a mistake to tie the epistemic state too tightly to any particular practical capacity. Could these capacities function as sufficient conditions for understanding? It may depend on the details. It seems possible that one person might be able to explain a concept they fail to understand by parroting the speech of another. But their capacity to explain does not in this case signal their understanding of the thing explained. One might reject that such an agent in fact can explain the relevant phenomena, but this seems like a linguistic dispute.

This leaves us with one last issue before I attempt to briefly reiterate the reasons why understanding is more promising than knowledge when it comes to responding to the Weak Guiding Intuition. This issue has to do with a potential complaint or objection based on the fit between the Weak Guiding Intuition and objectual understanding. On the one hand, we might worry that the value of understanding won't be distinctive, and so won't vindicate the Weak Guiding Intuition. On the other we might worry that understanding is too *recherche* to fill this important role. I think the correct response to these worries is to emphasize our native familiarity with objectual understanding. Though the term of art reference to it is unfamiliar, the state we are after is widespread and widely sought after. While we could construct a stipulative propositional attitude which is related to understanding but formally different in some small way (this task would be easier if we could get a better grip on understanding itself) we might have the makings of an objection.

If such a state could be produced then it might be argued that understanding does not have its value distinctively. But this objection seems to me perfectly general. We could stipulate a near-duplicate of any proposed epistemic state with some minor alteration. This isn't an objection to understanding filling this role but instead to any familiar doxastic attitude satisfying the Weak Guiding Intuition. According to this objector, the field of potential epistemic states is populous enough to preclude any one from standing out. But what this objection really seems to call out for is a clarification, and perhaps one that will help bridge the gap between my view and proponents of the view that the value of knowledge is exhaustive of the propositional attitude part of the full picture of epistemic value. It would be deeply unsatisfying to produce some unfamiliar, artificial state and hold that up as the doxastic attitude that represents special success from the epistemic point of view. I do not think this would be a fair characterization of my proposal. Understanding, though under-theorized, is far from unfamiliar.

To conclude, we've seen reason to think that insofar as we wanted to find out which doxastic attitude bore special epistemic value, we would need to distance ourselves from the instrumental value shared by most true belief and knowledge, among other epistemic states. Our emphasis on finding a state that has a unique value led us towards understanding, which we sketched in a way that made the features of that state that were inappropriate in the theory of knowledge, like a connection to reflection and offering reasons or explanations, stand out. It is just these features that make understanding a plausible candidate to fill this role. It is somewhat revisionary to deny that knowledge holds pride of place among the doxastic attitudes. But offering understanding instead allows a more satisfying account of success from the epistemic point of view while allowing us to maintain the kind of compact theory of knowledge that the last several decades of inquiry have converged on.

There is much more to do: more needs to be said about the nature of mere true belief. The role of junk knowledge in this debate could well be clarified. Perhaps a more direct response to the Instrumentalist about epistemic value is available. Certainly much more work needs to be done trying to get a clearer picture of objectual understanding. But I hope to have at least offered some reasons that support the idea that it is our understanding, not our knowledge, that represents our greatest epistemic achievement.

References

- Adam Carter, J. & Pritchard, Duncan (2015). Knowledge-How and Epistemic Value. *Australasian Journal of Philosophy* 93 (4):799-816.
- Allen, Julie (2003). G.e. Moore and the principle of organic unity. *Journal of Value Inquiry* 37 (3):329-339.
- Alston, William P. (1995). How to Think about Reliability. *Philosophical Topics* 23 (1):1-29.
- Alston, William (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.
- Audi, Robert (2009). Reliability as a virtue. *Philosophical Studies* 142 (1):43 - 54.
- Berker, Selim (2013). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review* 122 (3):337-393.
- Berker, Selim (2015). Reply to Goldman: Cutting Up the One to Save the Five in Epistemology. *Episteme* 12 (2):145-153.
- Bergström, Lars (1987). On the Value of Scientific Knowledge. *Grazer Philosophische Studien* 30:53-63.
- BonJour, L. (2010). The myth of knowledge. *Philosophical Perspectives* 24 (1):57-83.
- BonJour, L. (1990). Apriority and metajustification in BonJour structure of empirical knowledge-reply. *Philosophy and Phenomenological Research* 50 (4):779-782.
- Conee, Earl (1988). Why solve the Gettier problem? In D. F. Austin (ed.), *Philosophical Analysis*. Kluwer Academic Publishers. pp. 55--58.

- Conee, E. & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies* 89 (1):1-29.
- Dancy, Jonathan (2000). *Practical Reality*. Oxford University Press.
- DePaul, M. (2009). Does an Ugly Analysis Entail that the Target of the Analysis Lacks Value? In Pritchard, Haddock & Millar (eds.), *Epistemic Value*. Oxford: Oxford University Press. pp. 112--38.
- DePaul, Michael R. & Grimm, Stephen R. (2007). Review essay on Jonathan Kvanvig's the value of knowledge and the pursuit of understanding. *Philosophy and Phenomenological Research* 74 (2):498–514.
- Elgin, Catherine Z. (2009). "Is Understanding Factive?". In Adrian Haddock, Alan Millar & Duncan Pritchard (eds.), *Epistemic Value*. Oxford: Oxford University Press. pp. 322--30.
- Elgin, Catherine Z. (2004). True enough. *Philosophical Issues* 14 (1):113–131.
- Everson, Stephen. (Ed.) (1990) *Epistemology: Companions to Ancient Thought (Book 1)*. Cambridge University Press.
- Friedman, Jane (2018). Junk Beliefs and Interest-Driven Epistemology. *Philosophy and Phenomenological Research* 97 (3):568-583.
- Gardiner, Georgi (2017). Safety's Swamp: Against The Value of Modal Stability. *American Philosophical Quarterly* 54 (2):119-129.
- Goldman, Alvin (1979). "What is Justified Belief?". In George Pappas (ed.), *Justification and Knowledge*. Boston: D. Reidel. pp. 1-25.
- Goldman, Alvin I. (1993). Epistemic folkways and scientific epistemology. *Philosophical Issues* 3:271-285.

- Goldman, Alvin I. (2015). Reliabilism, veritism, and epistemic consequentialism. *Episteme* 12 (2):131-143.
- Goldman, Alvin I. & Olsson, Erik J. (2009). "Reliabilism and the Value of Knowledge". In A. Haddock, A. Millar & D. H. Pritchard (eds.), *Epistemic Value*. Oxford: Oxford University Press. pp. 19--41.
- Graham, Peter J. (2011). Does Justification Aim at Truth? *Canadian Journal of Philosophy* 41 (1):51-72.
- Greco, John (2007). The nature of ability and the purpose of knowledge. *Philosophical Issues* 17 (1):57–69.
- Grimm, Stephen R. (2006). Is understanding a species of knowledge? *British Journal for the Philosophy of Science* 57 (3):515-535.
- Grimm, Stephen R. (2012). The Value of Understanding. *Philosophy Compass* 7 (2):103-117.
- Grimm, Stephen R. (2016). The Value of Reflection. In Miguel Angel Fernandez (ed.), *Performance Epistemology: Foundations and Applications*. Oxford University Press.
- Grimm, Stephen R. (forthcoming). "Understanding and Transparency". In *Explaining Understanding: New Essays in Epistemology and the Philosophy of Science*. Routledge.
- Grimm, Stephen R. (2016). How Understanding People Differs from Understanding the Natural World. *Philosophical Issues* 26 (1):209-225.
- Grimm, Stephen R. (2014). "Understanding as Knowledge of Causes". In A. Fairweather (ed.), *Virtue Scientia. Virtue Epistemology and Philosophy of Science*. Springer.

Haddock, Adrian ; Millar, Alan & Pritchard, Duncan (eds.) (2009). *Epistemic Value*. Oxford University Press.

Hannon, Michael (2017). A solution to knowledge's threshold problem. *Philosophical Studies* 174 (3):607-629.

Harman, Gilbert (1986). *Change in View*. MIT Press.

Hills, Alison (2015). Understanding Why. *Noûs* 49 (2):661-688.

Horwich, P., In *The Aim of Belief*, Ed. Chan, T. (2013). Oxford.

Horwich, Paul (forthcoming). Is truth a normative concept? *Synthese*:1-12.

Horvath, Joachim (2016). Conceptual analysis and natural kinds: the case of knowledge. *Synthese* 193 (1):167-184.

Hyman, John (2010). The road to Larissa. *Ratio* 23 (4):393-414.

Jäger, Christoph (2011). Process Reliabilism and the Value Problem. *Theoria* 77 (3):201-213.

Jäger, Christoph (2011). Reliability and Future True Belief: Reply to Olsson and Jönsson. *Theoria* 77 (3):223-237.

Jäger, Christoph (2009). Why to believe weakly in weak knowledge: Goldman on knowledge as mere true belief. *Grazer Philosophische Studien* 79 (1):19-40.

Kaplan, Mark (1985). It's not what you know that counts. *Journal of Philosophy* 82 (7):350-363.

Khalifa, Kareem (2017). *Understanding, Explanation, and Scientific Knowledge*. Cambridge, UK: Cambridge University Press.

- Kim, Jaegwon (1994). Explanatory knowledge and metaphysical dependence. *Philosophical Issues* 5:51-69.
- Klein, Peter D. (1999). Human knowledge and the infinite regress of reasons. *Philosophical Perspectives* 13 (s13):297-325.
- Klein, Peter D. (2007). Human knowledge and the infinite progress of reasoning. *Philosophical Studies* 134 (1):1 - 17.
- Korsgaard, Christine M. (1983). Two distinctions in goodness. *Philosophical Review* 92 (2):169-195.
- Kvanvig, Jonathan (2009). The value of understanding. In Pritchard, Haddock & Millar (eds.), *Epistemic Value*. Oxford: Oxford University Press. pp. 95--112.
- Kyle, Brent G. (2013). Knowledge as a thick concept: explaining why the Gettier problem arises. *Philosophical Studies* 165 (1):1-27.
- Lewis, David (1996). Elusive knowledge. *Australasian Journal of Philosophy* 74 (4):549 – 567.
- Lipton, Peter (2009). Understanding without explanation. In H. W. de Regt, S. Leonelli & K. Eigner (eds.), *Scientific Understanding: Philosophical Perspectives*. University of Pittsburgh Press. pp. 43-63.
- Littlejohn, Clayton (2013). The Russellian Retreat. *Proceedings of the Aristotelian Society* 113 (3pt3):293-320.
- Maher, Patrick (1990). Acceptance without Belief. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1990:381-392.
- Olsson, Erik J. (2009). In defense of the conditional probability solution to the swamping problem. *Grazer Philosophische Studien* 79 (1):93-114.

- Olsson, Erik J. (2011). Reply to Kvanvig on the Swamping Problem. *Social Epistemology* 25 (2):173 - 182.
- Olsson, Erik J. (1999). Cohering with. *Erkenntnis* 50 (2-3):273 - 291.
- Olsson, Erik J. & Jönsson, Martin L. (2011). Kinds of Learning and the Likelihood of Future True Beliefs: Reply to Jäger on Reliabilism and the Value Problem. *Theoria* 77 (3):214-222.
- Piller, Christian (2009). Desiring the truth and nothing but the truth. *Noûs* 43 (2):193-213.
- Pritchard, Duncan (2009). Knowledge, Understanding and Epistemic Value. *Royal Institute of Philosophy Supplement* 64 (64):19-43.
- Pritchard, Duncan (2008). I—Duncan Pritchard: Radical Scepticism, Epistemic Luck, and Epistemic Value. *Aristotelian Society Supplementary Volume* 82 (1):19-41.
- Pritchard, Duncan (2007). Recent work on epistemic value. *American Philosophical Quarterly* 44 (2):85 - 110.
- Pritchard, Duncan (2005). Scepticism, epistemic luck, and epistemic angst. *Australasian Journal of Philosophy* 83 (2):185 – 205.
- Quine, Willard V. O. (1951). Two Dogmas of Empiricism. *Philosophical Review* 60 (1):20–43.
- Riggs, Wayne D. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research* 64 (1):79-96.
- Riggs, Wayne, Understanding, knowledge, and the *Meno* requirement, manuscript.

- Rinard, Susanna (2015). No Exception for Belief. *Philosophy and Phenomenological Research* 91 (2).
- Schechter, Joshua (2017). No Need for Excuses: Against Knowledge-First Epistemology and the Knowledge Norm of Assertion. In J. Adam Carter, Emma Gordon & Benjamin Jarvis (eds.), *Knowledge-First: Approaches in Epistemology and Mind*. Oxford University Press. pp. 132-159.
- Schwab, Whitney (2015). Explanation in the Epistemology of the Meno. *Oxford Studies in Ancient Philosophy* 48:1-36.
- Schwab, Whitney (2016). Understanding epistēmē in Plato's Republic. *Oxford Studies in Ancient Philosophy* 51:41-85.
- Setiya, Kieran (2013). Epistemic agency: Some doubts. *Philosophical Issues* 23 (1):179-198.
- Smith, Angela M. (1998). Knowledge and Expertise in the Early Platonic Dialogues. *Archiv für Geschichte der Philosophie* 80 (2):129-161.
- Sosa, Ernest (1985). Knowledge and Intellectual Virtue. *The Monist* 68 (2):226-245.
- Sosa, Ernest (1980). The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge. *Midwest Studies in Philosophy* 5 (1):3-26.
- Sutton, Jonathan (2005). Stick to what you know. *Noûs* 39 (3):359–396.
- Sylvan, Kurt (2018). Veritism Unswamped. *Mind* 127 (506):381-435.
- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.
- Weiner, Matthew (2009). Practical reasoning and the concept of knowledge. In Pritchard, Haddock & Millar (eds.), *Epistemic Value*. Oxford: Oxford University Press. pp. 163--182.

Wright, Crispin (1991). Scepticism and Dreaming: Imploding the Demon. *Mind* 100 (1):87-116.

Yamada, Masahiro (2011). Getting It Right By Accident. *Philosophy and Phenomenological Research* 83 (1):72-105.

Zagzebski, Linda (2003). Intellectual motivation and the good of truth. In Linda Zagzebski & Michael DePaul (eds.), *Intellectual Virtue: Perspectives From Ethics and Epistemology*. New York: Oxford University Press. pp. 135--154.

Zagzebski, Linda (2001). Recovering Understanding. In M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press.

Zagzebski, Linda (2003). The Search for the Source of Epistemic Good. *Metaphilosophy* 34 (1-2):12-28.

Zagzebski, Linda (1996). Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge. *Philosophical and Phenomenological Research* 60 (1):197-201.

Zimmerman, Michael J., "Intrinsic vs. Extrinsic Value", *The Stanford Encyclopedia of Philosophy* (Spring 2015 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2015/entries/value-intrinsic-extrinsic/>>.