

The Threshold Problem, the Cluster Account,
and the Significance of Knowledge

Abstract. The threshold problem is the task of adequately answering the question: “Where does the threshold lie between knowledge and lack thereof?” I start this paper by articulating two conditions for solving it. The first is that the threshold be neither too high nor too low; the second is that the threshold accommodate the significance of knowledge. In addition to explaining these conditions, I also argue that it is plausible that they can be met. Next, I argue that many popular accounts of knowledge cannot meet them. In particular, I lay out a number of problems that standard accounts of knowledge face in trying to meet these conditions. Finally, near the end of this paper, I argue that there is one sort of account that seems to evade these problems. This sort of account, which is called a cluster account of knowledge, says that knowledge is to be accounted for in terms of truth, belief and a cluster of epistemic properties and also that knowledge doesn’t require having all members of the cluster, but merely some subset.

I see someone walking towards me who looks like my friend. It’s probably her, but I don’t yet know; she’s still some distance away and difficult to make out. Gradually, as she comes closer and closer, she becomes easier and easier to see and I become more and more confident that she is indeed my friend. So by the time she is right in front of me, I have already known it was her for some time.

This yields a question: at what point did I come to know that the person walking towards me was my friend? And in general, where does the threshold lie between knowledge and lack thereof? The task of adequately answering this question is known as “the threshold problem.”¹

¹ This problem has been discussed in e.g. BonJour (2002: 43, 46, 48-9), BonJour (2003: 21-3), BonJour (2010b), Brown (2014), Grimm (2011), Hannon (2014), Hannon (2017), Hetherington (2002), Hetherington (2006).

Solving the threshold problem is surprisingly difficult. In particular, as I argue, it is quite difficult to specify a threshold that both accommodates the significance of knowledge and lies neither too high nor too low.

This paper has two main goals. First: many authors have tried to use the threshold problem to help illuminate the significance of knowledge. I wish to do so as well. In particular, on my version of the threshold problem, in order to adequately answer the question “Where does the threshold lie between knowledge and lack thereof?”, one must answer it in a way that accommodates the significance of knowledge. Because of this, in the course of evaluating proposed solutions to the problem, I will put forward and evaluate some accounts of the significance of knowledge.

Second, I wish to argue that most of the commonly-defended accounts of knowledge face major challenges in answering the threshold problem. This will serve to cast doubts on the plausibility of these accounts of knowledge, given that there is good reason to think the threshold problem can be solved. Much of my paper will be spent describing the challenges these accounts face. But near the end, I will also defend a positive conclusion. In particular, I will argue that a certain sort of account of knowledge, a sort that is rarely defended in contemporary epistemology, can avoid these challenges. The type of account I wish to defend is called a “cluster account of knowledge”. Roughly speaking, a cluster account of knowledge is an account of knowledge in terms of truth, belief, and a cluster of epistemic properties. On cluster accounts, knowledge does not require having every member of the cluster, but merely some subset.

This paper has three sections. The first clarifies my version of the threshold problem. The second argues that most of the commonly-defended accounts of knowledge face major

challenges in answering the threshold problem. The third discusses how cluster accounts can avoid these challenges and thereby solve the threshold problem.

1 Clarifying the threshold problem

As I noted above, solving the threshold problem is a matter of adequately answering the question “where does the threshold lie between knowledge and lack thereof?” In this section, I will discuss two conditions for adequately answering this question and also argue that it is plausible that they can be met.

1.1 The first condition: make the threshold neither too low nor too high

The first condition for solving the threshold problem is that the threshold should be neither too low nor too high. For instance, it should be neither so low that it grants me knowledge of my friend when I first spotted her, nor so high that it denies me knowledge of my friend when she was right in front of me.

How might we state this condition more precisely? I prefer to understand “too high” as implying some sort of skepticism² and “too low” as making knowledge too ubiquitous.³ Slightly more precisely: there are some propositions that are commonly taken to be known (e.g. that so-and-so is the president) and other propositions that are commonly taken to be unknown (e.g. that so-and-so will be the president in fifty years). By a skeptical view, I mean a view according to

² A similar idea is suggested at Hannon (2014: 1134) and implicit in Brown (2014: 185), Hetherington (2002: 124).

³ A similar idea is implicit in Brown (2014: 187).

which we fail to know many of the things that are commonly taken to be known.⁴ By a view that makes knowledge too ubiquitous, I mean a view according to which we in fact know many of the things that are commonly taken to be unknown.

I should note that being skeptical is a matter of degree. That is, some views are more skeptical than others. For example, a view on which we have no knowledge at all is extremely skeptical. By contrast, a view on which we have various forms of a priori and empirical knowledge but lack knowledge of inductive conclusions is still a skeptical view, but it is somewhat less skeptical. Because being skeptical is a matter of degree, meeting the first condition will be a matter of degree. That is, other things being equal, the more skeptical one's view is, the less well one meets the first condition. Similar things hold regarding making knowledge too ubiquitous. One further clarification: I'm not assuming that there's exactly one point between being too skeptical and making knowledge too ubiquitous; perhaps there's a range of acceptable intermediate views.

This way of understanding "too high" and "too low" makes it clear why it is an advantage of an account that it meets the first condition; we wish to avoid both radical skepticism and making knowledge far too ubiquitous.

I take it as fairly inarguable that it is desirable to avoid a highly skeptical position or make knowledge far too ubiquitous. But some have explicitly responded to the threshold problem by endorsing a threshold that yields an unacceptably skeptical position Bonjour

⁴ A terminological note: some may wish to hold that knowledge is an incoherent concept and thus that there is no knowledge. (See e.g. Weiner (2009).) I will consider this to be a form of skepticism, seeing it is a position on which there are many propositions we normally take ourselves to know but in fact do not. Thanks to Chris Buford for pressing me on what to say about someone who holds that knowledge is an incoherent concept.

(2010a). Even aside from the costs of endorsing an unacceptably skeptical position, there is a further problem that remains for such people. In particular, they have to acknowledge that in ordinary conversation, it sometimes seems appropriate to ascribe knowledge.

Because of this, a natural question arises, when is it appropriate to do so; where does the threshold lie between when it's appropriate to ascribe knowledge and when it isn't? And adequately answering this question requires offering a threshold that is neither too high nor too low. In short, even those who endorse skepticism should acknowledge that a puzzle remains to be solved, viz. a puzzle regarding when knowledge attributions are appropriate.⁵

1.2 The second condition: accommodate the significance of knowledge

The second condition is that the threshold has to accommodate the significance of knowledge.

As I shall make clear in a moment, I will be interpreting “significant” in a fairly liberal manner. But some have offered a version of the second condition that is less liberal than my own. In particular, some people suggest that knowledge is valuable in a way that lack of knowledge is not, and that an adequate answer to the threshold problem should enable us to explain it.⁶

Perhaps the most extreme version of this sort of statement comes from Lawrence BonJour, when he writes in his discussion of the threshold problem:

⁵ For others who note this, see e.g. Hannon (2014, 1134-5). Someone might have the following worry: whether or not an utterance is appropriate depends on conversational context and thus the appropriacy conditions for ascribing knowledge are context-sensitive. My response: this is consistent with what I say in this paper; I am making no assumptions regarding the truth or falsity of contextualism as an account of knowledge. Thanks to an anonymous reviewer for pressing me on this.

⁶ Note that in saying that the account should enable us to explain the significance of knowledge, I don't mean to be committing myself to the claim that every necessary condition on knowledge is significant; it is consistent with my view that there be three necessary conditions on knowledge and one of these be relatively insignificant, while the other two jointly explain knowledge's significance. Thanks to an anonymous associate editor for pressing me on this.

... knowledge ... is supposed to be a supremely valuable and desirable cognitive state, one whose possession marks the difference between full cognitive success and at least some degree of cognitive failure: knowledge is the epistemic summum bonum ... I will assume throughout that this is an absolutely essential condition for any adequate account of knowledge. BonJour (2010b: 58).⁷

One thing to note right away is that the view that BonJour presents here – the view that knowledge is the supremely valuable cognitive state – is deeply controversial. Indeed, even the more mild view that knowledge is always more valuable than lack thereof is rather controversial.⁸

Because of the controversy regarding these claims about the value of knowledge, I would like to avoid taking a stand on whether, and in what way, knowledge is valuable. Rather, as I mentioned above, I prefer to understand the second condition by talking about “significance” rather than “value”. In particular, I will understand the second condition as follows: an adequate answer to the threshold problem will enable an explanation of the significance of knowledge.⁹

⁷ See also Hannon (2014: 1120), Hannon (2017: 608).

⁸ For discussion, see e.g. Kvanvig (2003), Pritchard (2009).

⁹ This sort of understanding is related to that suggested by Lawrence BonJour when he writes, “Paradoxically enough, the weak conception, though it seeks to preserve the truth of common-sense attributions of knowledge, makes it impossible to give an account of why those attributions have any real epistemic significance or why they should matter very much for epistemology” BonJour (2010b: 62). It is also related to the widespread claim that the threshold should not be “ad-hoc” or “arbitrary.” BonJour (2010b: 68), Brown (2014: 179), Hannon (2014: 1125), Hannon (2017: 607), Hetherington (2006: 43).

To further clarify: I'm not going to assume that the significance of knowledge has something to do with the value of knowledge.¹⁰ But neither am I going to assume that the significance of knowledge does not have to do with the value of knowledge. Rather I'm going to remain neutral regarding what the significance of knowledge consists in.

The reason I am going to remain neutral is that I think that my arguments in this paper work for a large number of views regarding the significance of knowledge. So, for example, I think my argument works if you think that knowledge is the epistemic summum bonum. But I also think they work if you deny this, but nonetheless think that whether one has knowledge frequently matters to us in our day-to-day affairs.¹¹

Likewise, I'm going to remain neutral regarding the question: does knowledge matter in every circumstance? That is, it's consistent with my view that there might be certain cases in which it doesn't matter at all whether someone knows something. For example, perhaps it doesn't matter at all whether a certain person knows a certain piece of trivia.

I should note that significance is not an all or nothing affair; properties can be more or less significant. I wish to hold that, at the very least, knowledge has a high level of significance. I'm remaining neutral regarding whether it is of supreme significance, but that doesn't mean that I mean to be allowing that it is of minimal significance. As I noted above, I think this assumption

¹⁰ As I understand significance, something can be significant even if it is not ultimately valuable. For instance, suppose I'm supposed to meet a stranger at a bar, and I'm told that I will be able to recognize the stranger because he will be wearing an old fedora. In such a case, it is significant whether or not a given person at the bar is wearing an old fedora even if an old fedora is not valuable.

¹¹ I take it as fairly uncontroversial that knowledge frequently matters to us in our day-to-day affairs. But here is a piece of evidence: the word "know" is one of the most commonly used words in English; it's 47th according to one of the most comprehensive analysis of the language, right ahead of "will", "as", and "up". This is according to the "Corpus of Contemporary American English" Davies (2017) – the "Oxford English Corpus" puts "know" at 59 Corpus (2011). For comparison: the Corpus of Contemporary American English puts "moral" at 1467 Davies (2017).

is reasonable, given the prominent place of knowledge in our everyday lives and in epistemology.

I have now finished explicating my version of the threshold problem. That is, I am finished elaborating what I take to be two conditions on adequately answering the question “where does the threshold lie between knowledge and ignorance?”. As should be clear at this point, I’ve been at pains to avoid controversy in the statement of my conditions. In particular, all I have assumed is the following: it is desirable that we specify the threshold between knowledge and ignorance in a manner that (i) avoids skepticism, (ii) avoids making knowledge far too ubiquitous, and (iii) accommodates the significance of knowledge. But as we are about to see, despite the relative tameness of these conditions, it is quite difficult for commonly-defended accounts of knowledge to meet them.

2 Some problems facing commonly-defended accounts of knowledge in their attempts to solve the threshold problem

Now that I’ve introduced the threshold problem, I am almost ready to discuss some problems that face commonly-defended accounts of knowledge in their attempts to solve the threshold problem.

But before I begin stating these problems, I need to clarify some terms.

2.1 What I mean by “account of knowledge”

Let me start by discussing the term “account of knowledge”. A note about my goals: in the course of clarifying the term, I wish to argue for the following claim: there exists a true account of knowledge, as I use the term “account of knowledge”. The reason I wish to argue for this claim is that it will be useful later on. In particular, later on in this paper, I will be defending cluster accounts in part via an argument by elimination. That is, I will argue that there aren’t true accounts of knowledge of certain other forms and then conclude that there must exist a true account of knowledge of my favored form. This argument will only be persuasive insofar as it is antecedently plausible that there is a true account of knowledge.

By an account of knowledge I mean a non-trivial statement of the form: one has knowledge just in case such and such other conditions hold. Here’s an example: “One knows a proposition just in case one truly believes it and has enough evidence to appropriately assert it.”

By “non-trivial” here, I mean to exclude statements that reference knowledge among the conditions, either explicitly or via some term of art. So, for example, the following is a trivial account of knowledge: “one has knowledge just in case one has knowledge”. For another example: if one introduces the term of art “quasi-knowledge” as anything that, when combined with belief, yields knowledge, then the following is a trivial account of knowledge: “one knows a proposition just in case one believes it and quasi-knows it.”

I should emphasize that on my use of “account of knowledge”, the notion is very thin. To offer an account, one does not need provide a deep conceptual analysis. All one has to do is to offer conditions that are extensionally adequate.

There are at least three ways in which my notion of an account of knowledge is a thin one. First, it’s consistent with my usage that there can be multiple true accounts of knowledge.

All that it takes to offer a true account of knowledge is to correctly specify when one has knowledge and when one lacks it, and there may well be multiple ways to do so. Compare: as I use the term “account”, there are multiple true accounts of the term “equilateral triangle”. For example, one such account is that an equilateral triangle is a triangle with sides of equal length. Another is that an equilateral triangle is a triangle with angles of equal degree.

Second, as I use the term, there is no assumption of priority, grounding, etc. between knowledge and the “account” given of it. The sole requirement on something’s being a true account of knowledge is that it be extensionally adequate. So, as I am using the term “account” it’s possible to give a true account of A in terms of B and then give a true account of B in terms of A.

Third, as I use the term, even someone who denies that there are any simple straightforward principles linking knowledge and other properties can still endorse the view that there is a true account of knowledge. All one needs to accept in order to think that there is a true account of knowledge is to think that there’s some true principle of the form: one has knowledge just in case such and such conditions hold, but this principle can be very complex. This point will be highlighted later on – the sort of account of knowledge I favor – a cluster account – is indeed a very complicated account.

I said a moment ago that it is quite plausible that my first condition for solving the threshold problem can be met. Part of my case for this has been negative; I have not built too much into what is required to offer an account of knowledge. But I can also offer a positive argument for the claim that there is a true account of knowledge, as I have defined “account”.

Here is the argument: suppose you denied that there is a true account of knowledge. In other words, suppose you denied that there was an extensionally adequate statement of the form: one has knowledge just in case such and such conditions hold. Then you would be committed to denying that knowledge supervenes on other properties. But denying this is deeply problematic. In denying it, you thereby commit yourself to the claim that there can be two people who are exactly alike in every way with regards to every property besides knowledge, and yet be such that one has knowledge and the other doesn't. That is, in such a case, the two people would be equally reliable, have equal evidence, and so on. And yet one would have knowledge and the other not. This seems like a deeply unattractive result.

2.2 Clarifying some more terms

Next, there is the term “epistemic”. Roughly speaking, by epistemic components of an account of knowledge, I mean to pick out those properties besides truth and belief that figure in accounts of knowledge.¹² Examples that have been proposed include safety Sosa (1999), sensitivity Nozick (1981), reliability Goldman (1986), proper function Plantinga (1993), cognitive achievement Greco (2010), absence of luck Pritchard (2005), evidential support Feldman and Conee (1985), lacking false lemmas Armstrong (1973: 152), absence of defeaters Lehrer and Paxson (1969, and support of seemings Huemer (2001), among others.

It is worth emphasizing that in addition to these properties, my use of the term “epistemic” is also meant to include properties that make reference to practical considerations,

¹² This condition is somewhat rough. The reason is that some people think that knowledge does not require belief, but rather some similar but slightly different state, such as acceptance Myers-Schulz and Schwitzgebel (2013). So a little better is to say that epistemic properties are those properties – besides truth, belief, and other tightly related properties – that figure in an account of knowledge.

properties like being reliably believed enough to assert or evidentially supported enough to act on.¹³

Next, there is a distinction between “monist accounts of knowledge” and “pluralist accounts of knowledge”. In brief, monist accounts of knowledge (“monist accounts” for short) account for knowledge in terms of truth, belief, and a single epistemic property, while pluralist accounts of knowledge (“pluralist accounts” for short) account for knowledge in terms of truth, belief, and multiple epistemic properties. So, for example, the following is a monist account: knowledge is true belief that is sensitive. And the following is a pluralist account: knowledge is true belief that is both sensitive and lacks defeaters.¹⁴

Finally, there is a distinction between two sorts of epistemic properties: “pragmatist properties” are those related to our practical interests, they include properties like “having enough evidence to act on” and “being reliably believed enough to appropriately assert” and so

¹³ In light of this fact, I should acknowledge that different people use the term “epistemic” in different ways. For discussion, see e.g. Cohen (2016). Especially important to note is that some people use the term “epistemic” in a rather restrictive sense, so that properties that involve practical considerations do not count as “epistemic.” But I mean to be using the term more broadly than this. In short, I need a way to refer to properties besides truth and belief that appear in accounts of knowledge, and “epistemic” seems like a handy term to do so.

¹⁴ There are interesting questions about how to make this distinction more precise. For instance, suppose we defined a new conjunctive property, “sensitivity and lacking defeaters”, as follows: a belief has this conjunctive property just in case it is both sensitive and lacks defeaters. Suppose we offered an account on which knowledge was true belief that was sensitive and lacked defeaters. Is this a monist account or a pluralist account? My response: For the purposes of this paper, accounts of knowledge that define it in terms of such conjunctive properties should be considered pluralist accounts.

on. “Purist properties” are those that are not related to our practical interests; they include properties like reliability, evidential support, and so on.¹⁵

2.3 Problems for monist accounts of knowledge

I will start by raising some problems for monist accounts of knowledge.

First I shall discuss pragmatist monist accounts. As should be clear from the definitions above, a pragmatist monist account of knowledge is one that characterizes knowledge in terms of truth, belief, and a single epistemic property – in particular, an epistemic property that is connected to practical interests.

Pragmatist monist accounts face what I call “the diffusion problem.” In particular, while there are many pragmatist monist accounts, the properties they identify knowledge with, instead of clustering together, land at different places from one another.

For instance, one popular suggestion is that knowledge is closely tied to action.¹⁶ There are several different ways to spell this out, but let us take for illustration the claim that one knows that *p* just in case one can appropriately act as if *p*.

This property – whether it’s appropriate to act as if *p* – comes apart from the properties that other pragmatist monist accounts identify with knowledge. For instance, some suggest that

¹⁵ Likewise, purist accounts of knowledge are those that do not invoke any pragmatist epistemic properties, while pragmatist accounts of knowledge are those that invoke at least one pragmatist epistemic property. So, for example, the account on which knowledge requires true belief supported by one’s evidence is a purist account of knowledge, while the account on which knowledge requires true belief that one has evidence enough to act on is a pragmatist account of knowledge. For more on this distinction, see e.g. Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005).

¹⁶ For endorsement (sometimes tentative) of accounts of knowledge in terms of action, see e.g. Fantl and McGrath (2002), Fantl and McGrath (2009), Hawthorne (2004: 29-31), Hawthorne and Stanley (2008), and Stanley (2005).

one knows something just in case it is appropriate to assert it.¹⁷ But what it's appropriate to assert and what it's appropriate to act on can come apart.

One way they can come apart is when I'm acting in a relatively low stakes situation and talking with someone in a high stakes situation. So, for example, suppose I'm talking with someone about whether the bank will be open tomorrow while at the same time trying to decide whether I should go to the bank today or delay the trip until tomorrow. I might be in a good enough epistemic position with regards to the proposition that the bank will be open tomorrow that I can act on this proposition, say by delaying a trip to the bank until tomorrow. But perhaps the person I'm talking to is someone for whom it's extremely important that the bank be open, and in such a case my evidence that it'll be open won't be enough for it to be appropriate for me to outright assert that the bank will be open tomorrow.¹⁸

And what it's appropriate to assert and act on can come apart in the opposite direction; if I'm in a high stakes situation, but talking with someone in a relatively low-stakes situation, it might be inappropriate for me to act as if p but appropriate to assert p. So, for instance, if someone in a low stakes situation asks me if the bank will be open tomorrow, it might well be appropriate for me to assert that it will be, even if I can't act on this proposition, thanks to my own high-stakes situation.

Similar things hold for other properties often tied to knowledge, like the property that it's appropriate to stop inquiring as to the truth of a proposition or the property of flagging reliable

¹⁷ For endorsement (sometimes tentative) of this or related links between knowledge and assertion, see e.g. Adler (2002), Hawthorne (2004: 21-4), Williamson (2000: 238-70).

¹⁸ Similar cases about banks have appeared in a number of articles and books; one of the earliest versions appears at DeRose (1992: 913) and this particular ones I discuss are modeled on ones at Stanley (2005: 4-5).

informants.¹⁹ And it's not just me who thinks so; these accounts are regularly treated as rivals to each other.²⁰

Furthermore, each of these practical concerns can be tied to knowledge in a variety of ways. There is being reliable enough to act on, being evidentially supported enough to act on, and so on. And these come apart. For instance, one can have lots of evidence regarding something but not be particularly reliable with regards to it. (This might happen if one is bad at assessing one's evidence.) Likewise, one can be reliable regarding something, but lack evidence for it.²¹

In sum, while there are many pragmatist monist accounts, the properties they identify, instead of clustering together instead land at different places from one another. This causes problems regarding the condition that the threshold has to accommodate the significance of knowledge. In particular, because the properties picked out by monist accounts come apart from each other, acquiring one doesn't imply that one will have thereby acquired the others. And so even granting that a particular property picked out by a monist account is of some significance, it's dubious that it's of sufficient significance that we mark it using the term "knowledge" and focus so heavily on it.

¹⁹ For those offering accounts that tie knowledge to these properties, see e.g. Craig (1990), Hannon (2014), Kappel (2011), Kelp (2011).

²⁰ Indeed, the two bank examples I presented are used to illustrate such differences. See e.g. DeRose (2009), Stanley (2005: 5-6). And for some examples of people treating the property of its being appropriate to stop inquiring as to a proposition and the propriety of flagging reliable informants as rivals, see e.g. Kappel (2011), Kelp (2011), and Rysiew (2012). For a dissenting view, see Hannon (2015).

²¹ Here's an example of being reliable but lacking evidence: one could be a "naive chicken-sexer", someone who could reliably tell male baby chickens from female ones, while being unable to tell how one could do so. See e.g. Goldman (1975: 114-6), Pritchard (2005: 175-6).

Let me develop the point a little more fully. We have found two things in examining the various pragmatist epistemic properties. First, there are a large number of them. Second, by and large, for any given two, having one is neither necessary nor sufficient for having the other.

Here is a third thing to note: many of these properties are of some significance. After all, many of these pragmatist epistemic properties are such that accounts of knowledge have been offered in terms of them. And when people give accounts of knowledge in terms of some property, they tend to think that property is of some significance. In addition, many of the properties seem, on their face, to be significant – it seems to matter, at least a little, whether a belief is evidentially supported enough to act on, whether it is reliably believed enough to assert, whether it seems true enough to end one’s inquiry regarding it, and so on.

Given these three things: that there are a large number of epistemic properties, that they come apart, and that they are all of some significance, it seems hard to see how just one of these properties could be significant enough, on its own, to mark the difference between knowledge and ignorance. What could make that particular property so special, given that there are so many others out there, each of significance as well, and each coming apart from it?

In short, knowledge is highly significant. But if there are a large number of pragmatist epistemic properties, each of which is of some significance, and which comes apart from the others, then none, on its own, is highly significant. But if none is, on its own, highly significant, then in offering an account of knowledge in terms of one of them, one fails to accommodate the significance of knowledge.

One might object as follows: presumably, if one is a monist about pragmatic conditions on knowledge, one thinks that there are arguments for one’s favored condition that are

unsuccessful when it comes to other conditions. So one is not really going to be moved by the claim that there are other views that other philosophers – even those in the same broad camp as one’s self – accept.

My response: even if one thinks there are strong arguments for a certain pragmatist monist account, it seems to me that these arguments can be undercut by a failure for the account in question to accommodate the significance of knowledge. So, for example, suppose I support an account that characterizes knowledge in terms of having enough evidence to appropriately act. The reason I support this account is because I think it is relatively simple and avoids certain counterexamples that other accounts face. In offering these reasons, I haven’t yet offered any evidence that the property in question – having enough evidence to appropriately act – is especially significant, as compared with the other epistemic properties available. And thus, even though there are some good arguments in support of my account, it nonetheless comes with significant costs.²²

I just talked about pragmatist monist accounts, now let me turn to the other sort of monist account: purist monist accounts. Such accounts characterize knowledge in terms of truth, belief, and a single epistemic property. But this time, the epistemic property is not tied to practical interests.

The first thing to note is that even in adopting a purist monist account, one cannot entirely avoid the diffusion problem. As I illustrated earlier when I first defined the term “epistemic” there is a large variety of purist epistemic properties. And just as the pragmatist properties come apart, these purist ones do as well. For instance, as I noted, reliable belief

²² Thanks to an anonymous associate editor for pressing this point.

formation is neither necessary nor sufficient for evidential support. As with pragmatist epistemic properties: given the variety of purist epistemic properties and the fact that they come apart in this way, it is hard to see why one purist epistemic property would be sufficiently significant that knowledge would amount to having it.²³

In addition, almost all, if not all purist accounts face what I call “the sparseness problem”: it is hard to find a purist epistemic property that both accommodates the significance of knowledge and lies neither too high nor too low.²⁴ So, for example, take process reliabilism, according to which knowledge is true belief that is formed via a reliable process. Fifty percent reliability is significant but too low, a hundred percent reliability is significant but too high, and no level of reliability in between these is significant.

To illustrate further, suppose someone proposed that the threshold for knowledge was at 72.4 percent reliability. It’s hard to see why 72.4 percent reliability could be, in and of itself, particularly significant, as opposed to the other thresholds available.

It’s worth emphasizing that this isn’t a problem about vagueness. Even if we take a particular region, say the region between 71.3 and 75.9, it’s hard to see why this region could be

²³ One might think that defenders of purist accounts hold that their properties are significant in a way that other properties aren’t. So, for example, perhaps internalists think that internally accessible properties are important in certain ways that external properties are not. And thus, one might think that defenders of purist accounts avoid the diffusion problem.

My response: even if one thinks internal properties are significant, it’s hard to avoid the claim that externalist properties are too. Are not safety, sensitivity, reliability, etc. rather significant, albeit in different ways from internalist properties? But even setting this aside, it’s worth noting that there are many different sorts of internalist accounts. Is one an access internalist or a mental state internalist? If an access internalist, does one care about actual access or the ability to access? Is one an evidentialist? If so, what counts as evidence (seemings? phenomenal states? propositions one knows?), and what counts as evidential support (is it delivered by prima facie rules? is it a sui-generis relation? does it concern probabilistic connections?), etc. Internalists are strongly divided on these questions and it’s hard to believe that only one answer here is picking out a significant property. Thanks to an anonymous referee for pressing me on this point.

²⁴ Something like this problem is pressed at Grimm (2011: 719-21).

of particular significance, as compared with other regions between 50 percent and 100 percent. Why would we care so much about whether someone is above or below this particular region?

Compare baldness: even if it's vague at which point one becomes bald, there are certain regions that are more significant than others. For instance, the region between 0 and 10,000 hairs is more significant than the region between 170,000 and 180,000.

Similar things are true for other standard purist properties, such as evidential support; fifty percent likelihood given one's evidence is too low, a hundred percent too high, and there's no point or region that's particularly significant in between.²⁵ And so one way to make the case that the sparseness problem applies to every purist monist account is to run through purist monist accounts in turn.

Is there any more general fact that explains why the sparseness problem holds? It seems to arise from a fact about truth, that while being just as likely to be true as false is significant and while being true is significant, there isn't anything of significance in between. So purist accounts, which cannot invoke pragmatist properties, have no level of truth that is both intermediate and significant to latch onto.

2.4 Problems for moderately pluralist accounts of knowledge

Perhaps one might wish to avoid these problems. To do so, one will have to be a pluralist. There are two ways of being a pluralist, depending upon how many epistemic properties one's account invokes. In particular, some accounts invoke just a couple of epistemic properties, while others invoke more. I will use the term "moderate pluralist account" to refer to accounts of knowledge

²⁵ See e.g. Freedman (Forthcoming).

that account for knowledge in terms of just a couple of epistemic properties and “radical pluralist accounts” to refer to accounts of knowledge that account for it in terms of a greater number of epistemic properties.²⁶

As I see it, when it comes to the threshold problem, moderate pluralists face roughly the same problems that monists face. Let us start with pragmatist moderate pluralists, those that offer accounts in terms of properties, at least some of which are connected with our practical interests.

As with pragmatist monists, pragmatist moderate pluralists face the diffusion problem. Again, as a reminder, this problem arose in the case of pragmatist monists because there are a large number of epistemic properties, each of which came apart from the others. And it seems dubious that any of them is of sufficient significance that we mark it using knowledge. Switching to moderate pluralism from monism helps with the problem, but it doesn’t overcome it completely. In particular, while having two such properties is more significant than having just one, it still seems not quite significant enough to accommodate the significance of knowledge, given the large range of properties out there and the fact that they come apart.

In short, the problem is the following: given the broad array of pragmatist properties on offer, and given that they all come apart, it’s hard to see why any two of them would be of sufficient significance that they could allow us to mark the threshold between knowledge and ignorance.

²⁶ Many philosophers offer moderate pluralist accounts. In particular, many philosophers offer accounts of knowledge that include two epistemic properties, one to handle justification and the other to handle Gettier cases. For a large number of examples of such accounts, see Shope (1983). So, for example, some say that knowledge is true beliefs that is justified and lacks defeaters Shope (1983: 81-118). Such people count as offering moderate pluralist accounts. (Indeed, given this, there are arguably more moderate pluralists than there are monists. Thanks to an anonymous associate editor for pressing me on this point.) By contrast, there are relatively few philosophers who offer radical pluralist accounts. That said, there are some. For instance, David Lewis (arguably) offers one Lewis (1996).

Meanwhile, those purist moderate pluralists – those who characterize knowledge in terms of epistemic properties that do not mention practical interests – face the sparseness problem. As a reminder, this problem arose for purist monists because there are no significant purist properties that lie neither too high nor too low. This problem remains if we add in a second purist conditions. So, for example, I noted earlier that the threshold of 72.4 percent reliability isn't a significant threshold. This problem doesn't get solved if we switch to a threshold that combines two purist conditions as opposed to one. So, for example, the following is not a significant threshold: being at least 72.4 percent reliable and also having a likelihood of .628 conditional on one's evidence. If anything, it seems even less significant.

2.5 A problem for rigid radical pluralists

So far, I've enumerated some problems that arise for accounts of knowledge that are monist or moderately pluralist. So let me turn to looking at radical pluralist accounts – those that account for knowledge in terms of truth, belief, and a number of epistemic properties.

In discussing the different types of radical pluralist accounts, there is a final distinction I would like to make. It arises as a result of the following fact: when one gives an account of knowledge in terms of some epistemic properties, one has two options concerning the strictness of the connection. One can say each of these properties is necessary for knowledge, in which case one endorses what I call a rigid account, or one can deny this and say that at least some of the epistemic properties are not necessary for knowledge, in which case one endorses what I call a flexible account.

So, for example, take the account on which knowledge is true belief that is reliably formed and evidentially supported. This is a rigid account; it says that both of these epistemic properties are necessary for knowledge. By contrast, consider an account on which we specify a certain set of five epistemic properties and then say that knowledge requires having at least three out of this set of five. This is a flexible account because it does not say that each of these properties is necessary for knowledge. Rather, it merely requires three, but it doesn't matter which three.²⁷

In a moment, I will be discussing a problem that rigid radical pluralists face when it comes to solving the threshold problem. But before I offer this problem, I do want to note that rigid radical pluralists have some advantages over their less pluralist peers.

In particular, rigid radical pluralist accounts seemingly have a response to the diffusion problem, a problem that plagues less pluralist accounts. This problem, recall, was that, given the broad range of epistemic properties on offer, and given that they came apart from each other, it was hard to see how one or two could be sufficiently significant that we mark them using knowledge. Rigid radical pluralists seem to have an answer to this problem. They require that a number of these epistemic properties are possessed; this can be of great significance, even if each of these properties, on its own, is of limited significance.

Despite their advantages when it comes to the diffusion problem, rigid radical pluralists are not entirely home free. For there is a problem they face in trying to solve the threshold problem. I call the problem the “unhappy balance” problem.

²⁷ Most philosophers who offer accounts of knowledge offer rigid accounts of knowledge. Some exceptions, or at least possible exceptions, are Cohen (2004), Douven (2006), Gerken (2015), Hannon (2015).

To see the problem, first note the following fact: there are a number of epistemic properties that meet the following two conditions: (i) they have some amount of significance and (ii) saying they're necessary for knowledge seems to force us to adopt somewhat skeptical results in certain cases.

For instance, take the property of having enough evidence to act rationally on a proposition. This seems like a significant property and also seems tightly connected to knowledge. But if one claims that this property is necessary for knowledge, then one is pushed towards somewhat skeptical results.

To illustrate the sort of skeptical results I have in mind, we can consider Baron Reed's case involving jelly beans and electric shocks.²⁸ In this example, you are asked a series of questions. Each time you are asked a question, you can choose to answer it or to skip it. If you answer it correctly, you get a jelly bean, if you answer incorrectly, you get a very painful electric shock, and if you don't answer, nothing happens to you. Suppose you are asked if the battle of Hastings was fought in 1066. Suppose you are quite confident that it was and have fairly good evidence that it was; you distinctly remember learning this fact. But you know that if you're wrong, you'll get a painful shock. In such a case it seems irrational to say that the battle of Hastings was fought in 1066 – better to not answer at all, seeing as you'll get a painful shock if you're misremembering. So, if knowing requires having enough evidence to rationally act on a proposition, it follows that you don't know the battle was fought in 1066. But this seems unduly skeptical.

²⁸ It appears in Reed (2012); Jessica Brown has some similar cases that appear in Brown (2012) and Brown (2014).

The property of having enough evidence to rationally act is not the only epistemic property with this feature. As the history of attempts to analyze knowledge reveals, there are many epistemic properties with the following features: (i) they have at least some significance, (ii) they seem to be tightly tied to knowledge, and (iii) if we say they are necessary for knowledge, we are pressured towards somewhat skeptical results in some cases.²⁹

Taking on one such set of skeptical results is bad enough. But the problem is magnified for rigid radical pluralists, for they think that knowledge requires a number of epistemic properties. In short, rigid radical pluralists are in an unhappy situation: the more conditions they adds to knowledge, the better they are to accommodate its significance, but the more skeptical their accounts become.

I should emphasize that it is in part the rigidity of a radical pluralist's account that causes this problem. In other words, the problem is partly caused by the fact that a rigid radical pluralist thinks each of the epistemic properties in her accounts is necessary for knowledge. This means that each time a rigid radical pluralist adds an epistemic property to her account, the account becomes more skeptical. And this is because every time she adds a property to the account, there is yet another property that one must possess in order to have knowledge.³⁰

3 Flexible radical pluralists and the cluster view

²⁹ See e.g. Shope (1983) for a number of examples.

³⁰ Of course, this is assuming the claim – which I have already defended in discussing the diffusion problem – that the various epistemic properties are such that none is necessary nor sufficient for the others.

Let me pause and take stock. If what I have argued for earlier in this paper is correct, then two things are true.

First, my two conditions for solving the threshold problem can be met. After all, the conditions are relatively mild. They merely say, regarding the threshold supplied by a given account of knowledge, that (i) it should be neither too high nor too low (that is, it should avoid skepticism and avoid making knowledge far too ubiquitous), and (ii) it should accommodate the significance of knowledge.

Second, many standard accounts of knowledge face major problems in meeting these two conditions. I have discussed three such problems: the diffusion problem, the sparseness problem, and the unhappy balance problem. I argued that if one is a monist or moderate pluralist, one will face at least one of the first two problems. So to avoid the first two problems, one has to be a radical pluralist. I also argued that if one is a rigid radical pluralist, one will face the third problem. So to avoid the third problem while endorsing radical pluralism, one has to be a flexible radical pluralist.

When we combine these claims, what we get is some strong considerations in favor of flexible radical pluralist accounts of knowledge. In particular, such accounts avoid problems that every other type of account faces.

That said, the argument doesn't explain how flexible radical pluralist accounts can get around these problems and this would be nice to know. So, in this section, I would like to supplement this argument by discussing how flexible radical pluralist accounts of knowledge can solve the threshold problem in a way that avoids the problems that other accounts face.

I will focus on a certain type of flexible radically pluralist account that I think does a particularly nice job of doing so, namely the cluster account of knowledge.

The cluster account of knowledge, as its name suggests, is motivated by the following thought: there are a lot of epistemic properties that tend to cluster together. By this I mean that there are a large number of epistemic properties such that, typically, when one has one, one has the other. So, for example, typically, when one has enough evidence to act on a proposition, one also reliably believes it enough to appropriately assert it. Typically, when one safely believes a proposition, one sensitively believes it. Typically, when a proposition seems true to one, one lacks defeaters for it. And so on.

Motivated by this thought, cluster accounts of knowledge say that there is a cluster of epistemic properties and that knowledge requires having a number of properties in the cluster, but not all of them. So, for example, here is a toy model. A cluster account of knowledge might specify ten epistemic properties and say that knowledge requires truth, belief, and seven of these ten, any seven will do.

What, then, do cluster accounts say about the significance of knowledge? They say that knowledge is significant insofar as it is tied to a large number of epistemic properties, each of which has at least some significance. Part of the reason being tied to all of these properties is significant is the following: in saying that a person has knowledge, we communicate a great deal of information about them.³¹ So for example, we allow someone to defeasibly infer that the person has enough evidence to act on the proposition, reliably believes it enough to appropriately assert it, and so on for a number of other properties.

³¹ For a similar suggestion, see e.g. Cohen (2004: 487).

Cluster accounts can also help explain why the threshold is non-arbitrary. So long as the epistemic properties do not scatter randomly, but instead tend to cluster, we can use facts about where they cluster to help set a non-arbitrary threshold. One way of thinking about this is via analogy with multivariable regression. Suppose, for example, that we are trying to identify a measure for quality of life in a country. We might find multiple factors that relate to quality of life: factors related to money (average income, poverty level), factors related to health (average life expectancy, infant mortality rate), etc. While these factors all come apart, they tend to cluster together; if a country has a higher average income, it tends to have a higher life expectancy. And because of this we can, using statistical methods, try to create a single measure, quality of life, that best correlates with these.

One might, at this point, object as follows: given that the epistemic properties tend to cluster, couldn't monists appeal to the same explanation? So, for example, take the property of having enough evidence to act on a proposition. Suppose that I'm correct that, typically, when one has this property, one believes it reliably enough to assert it and vice versa. In that case, being told that someone has enough evidence to act on a proposition defeasibly gives one information regarding whether they have enough information to reliably believe it. And likewise, for the other factors. So it might appear that monist accounts can appeal to the same account of the significance of knowledge as the cluster view can.

My response: such accounts will be less explanatory than cluster accounts for two closely related reasons. First, while there will be such a link between a certain account's favored epistemic property and the other epistemic properties, it will be less tight than the link is on the cluster account between the cluster-property and the other properties. In other words, call one's

avored epistemic property F and call C the cluster property. While possessing F may provide some defeasible evidence that one possesses a given epistemic property, E, and vice versa, it does not provide as much evidence as one's possessing C does.

To see this point, it's helpful to consider an analogy. Suppose I'm creating a television program to be watched by people across the United States. Each day, I will present a brief weather report. During this report, I will only have time to report a single temperature reading. I can either report the temperature for a given location (say Santa Fe, New Mexico) or report the average temperature for the country as a whole. Reporting the temperature for Santa Fe each day will give my viewers some information, no matter where they live. After all, the temperature in Santa Fe, New Mexico is correlated with the temperature in other places. But, by and large, reporting the average temperature for the country will give the viewers even more information. This is because there will be some local fluctuations which affect the region around Santa Fe but do not greatly affect other locations. Offering the average temperature for the country as a whole allows me to screen out these fluctuations and thus provide more useful information to viewers.

Second, in certain cases, a given epistemic property will come apart from most of the other epistemic properties. In these cases, the fact that one has the given epistemic property does not provide defeasible information that one has the other epistemic properties.

To illustrate, take the property of sensitivity. By definition, one sensitively believes a proposition just in case, if it was false, one wouldn't believe it. In some cases, this property comes apart from most of the other epistemic properties. For example, suppose I am a property manager running an apartment complex with a strict no-animal policy. Walking past a tenant's apartment, I see his pet labradoodle running around his living room and come to believe that he

has an animal in his apartment. Suppose furthermore there been no animal in his apartment (e.g. because he took his dog out for a walk), I would have seen a cat-shaped pillow in his living room and still believed that he had an animal in his apartment. In such a case, my belief that he has an animal in his apartment is not sensitive; I would have believe that he had an animal in his apartment even if he hadn't. But this belief does have most of the other standard epistemic properties; it's justified, I have good evidence for it, I lack defeaters for it, and so on. Because sensitivity in this case comes apart from most of the other epistemic properties, facts regarding whether my belief is sensitive provide little defeasible information indeed regarding the other epistemic properties.

How does the cluster model avoid the problems that faced the earlier accounts? First, there was the diffusion problem. This problem arose because while there are a lot of different epistemic properties, each of the properties comes apart from the others, and thus it is hard to see how an account of knowledge in terms of a mere one or two of these could accommodate the significance of knowledge. Someone who endorses the cluster account can avoid this problem – she can tie knowledge to a large number of other properties and thus, as we have just seen, accommodate its significance, even while acknowledging that these properties come apart.

Second, there is the sparseness problem. This problem concerned purist properties, those that are divorced from our practical interests. It arose because there seem to be few, if any, purist properties that are significant and that also lie neither too low nor too high.

Here, someone offering the cluster account has several options. One option is to endorse pragmatism and say that some of the properties in terms of which knowledge is to be accounted for are pragmatist properties. The cluster theorist can thereby avoid the sparseness problem,

which only arises for purist accounts. Of course, in doing so, she would take on the burden of endorsing pragmatism, which is a controversial view.³²

A second option is to offer a purist cluster account. In particular, one can argue that there is a set of purist properties that tend to cluster together, in the sense that if you have one, you tend to have the others. And then one can offer a purist cluster account of knowledge terms of this set.

For illustration, let me sketch out such an account. Some say that knowledge requires evidential support and understand evidential support in terms of being the best explanation of one's evidence.³³ And, furthermore, it is common to understand the notion of best explanation in terms of theoretical virtues like simplicity, explanatory power, coherence, and so on. This way of understanding best explanation seems to be a cluster view; it is both radically pluralist (there are a lot of properties connected with being the best explanation) and flexible (none is required for something to be the best explanation, e.g. lack of simplicity can be made up for by a great amount of explanatory power, coherence, etc.). And so in endorsing these views, one could

³² One might object as follows: Imagine an account on which you know that P iff you truly believe it and two out of a certain set of four pragmatist conditions. Suppose we agree that a threshold of .628 probability is not significant. It seems that in exactly the same way, meeting two of the four conditions is not significant. (Why not one? three? four?)

There are a number of responses available, but for the sake of brevity, I'll quickly sketch out one, which I find particularly attractive. Perhaps the number of conditions required for knowledge is selected so as to maximize informativeness. If we required all four conditions, we would learn a lot whenever we heard that someone had knowledge; we'd know that all four conditions were satisfied. That said, we'd rarely have the opportunity to truly ascribe knowledge, because having knowledge would require meeting all four conditions, which is rather demanding. By contrast, if we require merely one of the conditions, then we could ascribe knowledge fairly often, but learning that someone knows something would be not particularly informative, because it would merely indicate that one of the four conditions was met. The best account of knowledge would be one that would maximize informativeness; it would be one that could be ascribed somewhat frequently and at the same time, when ascribed, provide a fair bit of information. Thanks to an anonymous referee for pressing me on this point.

³³ See e.g. McCain (2011).

endorse a purist cluster account of knowledge and get around the sparseness problem. In particular, one would understand knowledge in terms of a cluster of purist properties, properties like simplicity, explanatory coherence, and so on.³⁴

Third, there was the unhappy balance problem. This problem arose because there was a general tension between tying knowledge to too few properties, and thus being unable to accommodate the significance of knowledge, and tying it to too many, and thus being forced into skepticism.

Cluster accounts avoid this problem. They do this by remaining flexible. This means that they can tie a large number of epistemic properties to knowledge without thereby deriving unduly skeptical results; after all, they don't say that the properties are strictly necessary for knowledge, but rather that the tie is somewhat weaker.

³⁴ One might worry that this sort of account faces the diffusion problem. After all, it seems to exclude a number of epistemic properties that have been thought important for knowledge; arguably reliability, sensitivity, and safety are not part of being the best explanation.

My response: I originally presented the diffusion problem in the context of discussing monist pragmatist accounts. There I pointed out that there are a large number of monist pragmatist epistemic properties and that each seems to come apart from the others. I questioned whether any, on its own, could accommodate the significance of knowledge.

This sort of objection doesn't seem, at least not straightforwardly, to apply to best-explanation-style account I just sketched. Even if the individual properties that figure in the best-explanation-style account have moderate significance, when we combine them, it is possible to derive something of stronger significance.

Perhaps the thought is the following: even if we grant that the property of being the best explanation is significant, if we could add in further properties that were left out of the notion of being the best explanation, and then we'd find something of even more significance.

But this is not clear to me. Perhaps, in combining the best explanation with further properties, we'd end up with something that was so weakly attached to these properties that it ended up not being particularly significant. Compare: the property of being happy is significant and the property of being wealthy is significant, but that doesn't mean that, in creating some property combining happiness and wealth, we'd end up with something particularly significant. That is, perhaps happiness and wealth are sufficiently different in key ways that a property combining happiness and wealth wouldn't, on its own, be particularly significant, or at least not as significant as the property of happiness is on its own. Thanks to an anonymous referee for pressing me on this point.

In sum, then, as we have seen, cluster accounts of knowledge seem to avoid the three problems faced by other sorts of accounts in their attempt to solve the threshold problem.

4 Conclusion

Let me summarize the results of this paper. If the arguments in this paper are correct, it appears that cluster accounts of knowledge can, unlike other accounts of knowledge, meet my two conditions for solving the threshold problem. And there is good reason to think the two conditions can be met. This strikes me as a relatively strong argument in favor of a view that almost no one endorses.

Of course, I cannot pretend to have decisively demonstrated that no other account can possibly solve the threshold problem. Nor have I examined the various criticisms one might have of the cluster account of knowledge, and whether these outweigh its advantages when it comes to solving the threshold problem. And finally, I have not endorsed a particular cluster account of knowledge. And because I haven't endorsed a particular cluster account of knowledge, I haven't said exactly where I personally think the threshold lies; different cluster accounts will give different answers to this question.

Nonetheless, I hope, at the very least, to have given some reason to think that we should start taking the cluster account of knowledge more seriously.

I opened this paper by telling a story and then posing a question about it. I would like to close this paper by returning to the story and answering the question.

My story involved my friend walking towards me and asked when it was that I came to know that it was my friend, as opposed to someone else. Here is what the cluster model would

say about this question. It would say that as my friend walks towards me, I gradually am put in a better and better epistemic position. For a while, I don't acquire many new epistemic properties. Then, at some point, I start to acquire a bunch in quick succession. For example, perhaps at some point midway through her walk I become able to rationally act on the proposition (by waving). Maybe shortly thereafter I have enough evidence to assert the proposition (by saying to my companion: "Look, there's Jo"). And perhaps shortly after that, I lose all defeaters for the proposition. And it is around this region, this region at which I am acquiring epistemic properties in quick succession, at which I came to know it was my friend. This, then, is where there threshold between knowledge and ignorance lies.³⁵

References

- Adler, J. E. 2002. *Belief's Own Ethics*. MIT Press.
- Armstrong, D. M. 1973. *Belief, Truth, and Knowledge*. Cambridge University Press.
- BonJour, L. 2002. *Epistemology*. Rowman and Littlefield.
- BonJour, L. 2003. 'The regress problem and foundationalism.' In *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Blackwell.
- BonJour, L. 2010a. 'Externalism/internalism.' In Dancy, J., Sosa, E., and Steup, M., editors, *A companion to epistemology*, pages 364–369. Blackwell, 2nd edition.
- BonJour, L. 2010b. 'The myth of knowledge.' *Philosophical Perspectives*, 24:57–83.
- Brown, J. 2012. 'Practical reasoning, decision theory, and anti-intellectualism.' *Episteme*, 9(1):43–62.
- Brown, J. 2014. 'Impurism, practical reasoning, and the threshold problem.' *Noûs*, 48(1):179–92.
- Cohen, S. 2004. 'Knowledge, assertion, and practical reasoning.' *Philosophical Issues*, 14(1):482–91.
- Cohen, S. 2016. 'Theorizing about the epistemic.' *Inquiry*, 59(7-8):839–57.
- Corpus, O. 2011. 'The OEC: Facts about the language.'
- Craig, E. 1990. *Knowledge and the State of Nature*. Oxford University Press.
- Davies, M. 2017. 'Word frequency data.'

³⁵ For helpful comments, I'm grateful to Chris Buford, Michael Hannon, and audiences at Cyncorp and the Pacific APA.

- DeRose, K. 1992. 'Contextualism and knowledge attributions.' *Philosophy and Phenomenological Research*, 52(4):913–929.
- DeRose, K. 2009. *The case for contextualism*. Oxford University Press.
- Douven, I. 2006. 'Assertion, knowledge, and rational credibility.' *Philosophical Review*, 115(4):449–85.
- Fantl, J. and McGrath, M. 2002. 'Evidence, pragmatics and justification.' *The Philosophical Review*, 111(1):67–94.
- Fantl, J. and McGrath, M. 2009. *Knowledge in an uncertain world*. Oxford University Press.
- Feldman, R. and Conee, E. 1985. 'Evidentialism.' *Philosophical Studies*, 48:15–34.
- Freedman, K. L. Forthcoming. Quasi-evidentialism: Interests, justification, and epistemic virtue. *Episteme*.
- Gerken, M. 2015. 'The roles of knowledge ascriptions in epistemic assessment.' *European Journal of Philosophy*, 23(1):141–61.
- Goldman, A. 1975. 'Innate knowledge.' In Stich, S. P., D., *Innate Ideas*, pp. 111–120. University of California Press.
- Goldman, A. 1986. 'What is justified belief.' In Pappas, G., ed., *Justification and Knowledge*. Reidel.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge University Press.
- Grimm, S. R. 2011. 'On intellectualism in epistemology.' *Mind*, 120(479):705–733.
- Hannon, M. 2014. 'Fallibilism and the value of knowledge.' *Synthese*, 191:1119–1146.
- Hannon, M. 2015. 'The importance of knowledge attributions.' *Philosophy Compass*, 10(12):856–66.
- Hannon, M. 2017. 'A solution to knowledge's threshold problem.' *Philosophical Studies*, 174:607–29.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Hawthorne, J. and Stanley, J. 2008. 'Knowledge and action.' *The Journal of Philosophy*, 105(10):571–590.
- Hetherington, S. 2002. *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*. Oxford University Press.
- Hetherington, S. 2006. 'Knowledge's boundary problem.' *Synthese*, 150:41–56.
- Huemer, M. 2001. *Skepticism and the Veil of Perception*. Rowman and Littlefield.
- Kappel, K. 2011. 'On saying that someone knows: themes from Craig.' In Haddock, A., Millar, A., and Pritchard, D. H., ed., *Social Epistemology*. Oxford University Press.
- Kelp, C. 2011. 'What's the point of "knowledge" anyway?' *Episteme*, 8(1):53–66.
- Kvanvig, J. L. 2003. *The value of knowledge and the pursuit of understanding*. Cambridge University Press.
- Lehrer, K. and Paxson, T. 1969. 'Knowledge: undefeated justified true belief.' *The Journal of Philosophy*, 66:225–37.
- Lewis, D. 1996. 'Elusive knowledge.' *Australasian Journal of Philosophy*, 74(4):549–567.
- McCain, K. 2011. *Inference to the best explanation and the external world: a defense of the explanationist response to skepticism*. PhD thesis, University of Rochester, Rochester, New York.

- Myers-Schulz, B. and Schwitzgebel, E. 2013. 'Knowing that p without believing that p.' *Noûs*, 47(2):371–84.
- Nozick, R. 1981. *Philosophical Explanations*. Harvard University Press.
- Plantinga, A. 1993. *Warrant and Proper Function*. Oxford University Press.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford University Press.
- Pritchard, D. 2009. 'Knowledge, understanding and epistemic value.' In *Epistemology*, Royal Institute of Philosophy supplements, volume 64, pp. 19–43. Cambridge University Press.
- Reed, B. 2012. 'Resisting encroachment.' *Philosophy and Phenomenological Research*, 85(2):465–72.
- Rysiew, P. 2012. 'Epistemic scorekeeping.' In Brown, J. and Gerken, M., ed., *Knowledge Ascriptions*, pages 270–93. Oxford University Press.
- Shope, R. K. 1983. *The Analysis of Knowing: A Decade of Research*. Princeton University Press.
- Sosa, E. 1999. 'How to defeat opposition to Moore.' *Philosophical Perspectives*, 13:141–153.
- Stanley, J. 2005. *Knowledge and practical interests*. Clarendon Press.
- Weiner, M. 2009. 'The (mostly harmless) inconsistency of knowledge ascriptions.' *Philosophers' Imprint*, 9(1):1–25.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford University Press.