Episteme (2026), 1–15
doi:10.1017/epi.2026.10137

episteme

ARTICLE

Metaepistemology and the Value Problem

Sophia Dandelet

Philosophy, University of Cambridge, UK
Email: dandelet@wisc.edu

(Received 7 June 2025 UTC; revised 12 May 2026 UTC; accepted 23 July 2026 UTC)

Abstract

What makes knowledge more valuable than mere true belief? This is the so-called *value problem*. Some epistemologists think that providing a solution to this problem is a key desideratum for theories of knowledge. One simple argument for this position starts with the claim that knowledge is intuitively more valuable than true belief. A more subtle and ambitious argument is rooted in the idea that theorists about knowledge must explain why knowledge is distinctively valuable, on pain of embarrassment; for without such an explanation, the thinking seems to go, epistemologists must confront the awkward possibility that they are spending time and energy on an unworthy subject. In what follows, I flesh out each of the arguments sketched above in a number of ways, and I argue that none of these versions succeed. The moral is that the value problem does not have the metaepistemological significance that is sometimes attributed to it.

Keywords: Value problem; metaepistemology; value of knowledge; function-first epistemology

What makes knowledge more valuable than mere true belief? This is the so-called *value problem*.

Some epistemologists think that providing a solution to this problem is a key desideratum for theories of knowledge. One simple argument for this position starts with the claim that knowledge is intuitively more valuable than true belief, and then adds the familiar metatheoretical premise that a theory of X should account for common intuitions about X. A more subtle and ambitious argument is rooted in the idea that theorists about knowledge must explain why knowledge is distinctively valuable, on pain of embarrassment; for without such an explanation, the thinking seems to go, epistemologists must confront the awkward possibility that they are spending time and energy on an unworthy subject.

In what follows, I flesh out each of the arguments sketched above in a number of ways, and I argue that none of these versions succeed. The moral is that the value

Thanks to audience members at the 2025 Joint Session of the Aristotelian Society and the Mind Association for questions and comments that helped me improve the manuscript.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

problem does not have the metaepistemological significance that is sometimes attributed to it. This is not to say that we should stop trying to solve the problem. It is to say, rather, that we should not expect our search for a solution to push us towards certain theories of knowledge and away from others.

1. Arguments from intuition

Sometimes epistemologists appeal to intuition to argue that the value problem should be treated as a starting point for theorizing about knowledge. They invoke a purportedly widespread, pre-theoretical intuition to the effect that knowledge is more valuable than mere true belief – call this the *value intuition* – and then they claim that our theory of knowledge should somehow accommodate this intuition.

Different versions of the argument from intuition are distinguished by different conceptions of what it takes to accommodate an intuition. Here are eight options:

What is it to accommodate an intuition? Eight options.
A theory accommodates an intuition that p just in case:

1. it explains, on its own, why we have the intuition that p.
2. it explains, in conjunction with some other facts, why we have the intuition that p.
3. it does not preclude a good explanation of why we have the intuition that p, i.e., it does not render mysterious the fact that we have this intuition.
4. it is compatible with the fact that we have the intuition that p.
5. it explains, on its own, why p is true.
6. it explains, in conjunction with some other facts, why p is true.
7. it does not preclude a good explanation of why p, i.e., it does not render p mysterious.
8. it is compatible with p.

Notice that possibilities (1–4) and (5–8) form two natural groups. (1–4) offer accounts of intuition accommodation according to which a theory might accommodate an intuition even while implying that the intuition is misleading. I'll call this kind of accommodation *error-allowed accommodation*. (5–8), by contrast, offer accounts of intuition accommodation according to which accommodating an intuition requires committing to the correctness of that intuition. Let's call this kind of accommodation *vindicatory accommodation*. Within each of these genera are more ambitious and less ambitious species. The most ambitious ones (1 and 5) construe accommodation as complete explanation of the relevant fact – the fact that we have some intuition or the fact that the intuition is true, depending on whether we are in the error-allowed or vindicatory branch of the accommodation tree. The most modest ones (4 and 8) construe accommodation as mere compatibility with the relevant fact. And between the two poles, we have pictures of accommodation as partial explanation of the relevant fact (2 and 6) and as compatibility with a satisfying explanation of the relevant fact (3 and 7).

With this taxonomy of intuition accommodation in place, we can start to examine a variety of ways in which the value intuition might be used to generate constraints on theories of knowledge.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

1.1. The value intuition and error-allowed accommodation

Let's begin by considering constraints that ask for error-allowed accommodation of the value intuition. Pritchard (2009: 19), for instance, suggests that a theory of knowledge should explain why we have the intuition that knowledge is distinctively valuable:

> It is a widespread pre-theoretical intuition that knowledge is distinctively valuable . . . Given this fact, however, it is obviously important to a theory of knowledge that it is able to offer a good explanation of this fact.

Pritchard goes on to point out that there is more than one way of explaining the existence of an intuition: We could vindicate the intuition, i.e., say what makes it correct. Alternatively, we could explain the intuition away, i.e., say why we have the intuition even though it is incorrect. So, it is clear that Pritchard is asking for error-allowed accommodation of the value intuition.[1]

On the most ambitious interpretation of this demand, it calls for a theory of knowledge that completely explains why we intuit that knowledge is more valuable than mere true belief; a slightly less ambitious interpretation calls for a theory of knowledge that at least partly explains why we have the value intuition. Both of these interpretations invoke a general metatheoretical premise about what philosophical theories can and should do: namely, that if p is a widespread, pre-theoretical intuition about X, it is obviously "important to a theory" of X that it plays some role in explaining why we have the intuition that p. It isn't obvious, however, that wholly providing or even just contributing to such an explanation is always the job of a philosophical theory of X, particularly in cases where the intuition in question is mistaken. A more plausible hypothesis is that explaining mistaken intuitions about X is at least sometimes a job for psychology, history, or anthropology, rather than a job for philosophy of X.

If you are moved by this criticism, you might be tempted by a more modest version of Pritchard's demand. For instance, you might ask for a theory of knowledge that is simply compatible with a good explanation of why we have the intuition that knowledge is distinctively valuable. But once we have acknowledged that explaining the existence of the value intuition might fall entirely to a non-philosophical discipline, this constraint seems to cross the line between modesty and meaninglessness – for it is difficult to see how any philosophical theory of knowledge could preclude the possibility of giving an excellent psychological, historical, or anthropological explanation of why people have a certain intuition about knowledge.

The same problem arises even more acutely for the most modest version of a demand for error-allowed accommodation of the value intuition, which asks for a theory merely compatible with the fact that we have the value intuition. I can't think of any existing theory of knowledge that falls short of this desideratum, presumably because there are no theoretical gains to be had by positing that there definitely is *not* a widespread, pre-theoretical intuition to the effect that knowledge is distinctively valuable.

It seems, then, that demands for error-allowed accommodation of the value intuition are caught between a rock and a hard place. Either they exaggerate the proper role of philosophy or they fail to generate meaningful constraints on theories of knowledge.

[1]Meylan (2007) also indicates, in footnote 2 on page 85, that she is interested in error-allowed accommodation of the value intuition.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

Notice that the concerns just sketched do not arise because of any special features of the value intuition *per se*. Versions of these concerns arise for many philosophical theories that are motivated by a demand for error-allowed accommodation. In a nutshell, this is because explaining why people have a mistaken intuition is not always a job, either in whole or in part, for philosophy. So, if we claim to be perfectly happy with a theory of *X* that allows for some widespread intuition to be mistaken, then it may not make sense for us to demand a theory of *X* that explains, either in whole or in part, why we have the intuition in question – that is, it may not make sense for us to demand one of the more ambitious varieties of error-allowed accommodation. And, while it remains clear why we should want one or both of the more modest varieties of error-allowed accommodation, this desideratum usually won't provide a very helpful theoretical constraint – because a philosophical theory of *X* is (in most cases) not the kind of thing that could conflict with a good empirical explanation of why we have a certain mistaken intuition about *X*, let alone the bare fact that we have this intuition.[2]

In many cases, then, it will be better to frame desiderata motivated by intuition in terms of vindicatory accommodation rather than error-allowed accommodation. As we are about to see, however, this move will not help matters when it comes to the value intuition.

1.2. The value intuition and vindicatory accommodation

Recall: A theory accommodates the intuition that *p* in a vindicatory way only if the theory does not imply that *p* is false. There are more and less ambitious versions of this species of accommodation. The most ambitious versions ask for complete or partial explanation of why *p*. The less ambitious versions ask for mere compatibility with a good explanation of why *p*, or mere compatibility with *p* itself.

Jones and Zagzebski seem to ask for vindicatory accommodation of the value intuition:

> All this indicates not only a conviction that knowing is different from merely true believing, but also that it holds some value over and above merely true believing. Given this, it would seem reasonable to demand from a theory of knowledge that it be able to explain why knowledge holds this value. (Jones 1997: 424)

> Philosophers have traditionally regarded knowledge as a highly valuable epistemic state, perhaps even one of the great goods of life. At a minimum, it is thought to be more valuable than true belief. Contemporary proposals on the nature of knowledge, however, make [this] difficult to understand . . . Some of the most common theories cannot even explain why knowledge is better than true belief. (Zagzebski 2003: 12)

Jones' passage is unambiguous: he is clearly asking for one of the more ambitious varieties of vindicatory accommodation, since he says that a theory of knowledge should be able to *explain* why knowledge has some value over and above that of mere true belief. Zagzebski's

[2]I have expressed the general version of my worry about error-allowed accommodation in "hedgy" language – *usually*, *in most cases* – because I think that there are certain philosophical projects to which this worry does not apply. These notably include philosophical projects that (a) are highly intuition-driven and (b) overlap substantially with empirical theories, like certain projects in formal semantics.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

passage, by contrast, allows a few interpretations. If we focus on the penultimate sentence – "Contemporary proposals on the nature of knowledge, however, make [this] *difficult to understand*" – Zagzebski seems to be asking for one of the more modest varieties of vindicatory accommodation, i.e., compatibility with a good explanation of why knowledge is better than true belief or even just the (purported) fact that knowledge is better than true belief. But if we focus on the last sentence – "Some of the most common theories cannot even explain why knowledge is better than true belief" – Zagzebski seems to be asking for one of the more ambitious varieties of vindicatory accommodation, i.e., complete or partial explanation of why knowledge is better than true belief.

1.2.1. More ambitious vindicatory accommodation

Let's consider each of these desiderata, starting with the more ambitious ones. Should we demand a theory of knowledge that explains, either on its own or in part, why knowledge is more valuable than mere true belief? I don't think so.

An analogy will help me explain why not. Suppose, as seems quite likely, that there is a widespread, pre-theoretical intuition to the effect that water is valuable. Now consider three philosophical theories about what water is. One of these theories explains on its own why water is valuable, while another only provides an explanation of this fact when it is paired with certain facts that are not part of the theory. The third theory plays no role at all in explaining why water is valuable; the explanation is provided wholly by facts external to the theory.

If this is all we know about these three theories of water, do we have any reason to prefer the first theory to the second, or the first and second to the third? It seems to me that we do not. If it is indeed true that water is valuable, then this fact must have some explanation; we can all agree on that. But there is no reason to suppose at the outset that this explanation must be the purview, either in whole or in part, of a philosophical theory of water. For all we have said so far, the explanation could come from other quarters.

Now, matters are different if what we want to explain is not simply why water is valuable, but rather, why water is *essentially*, *necessarily*, or *finally* valuable. Explanations of such facts *are* the special purview of philosophy. But, notice: while it is quite plausible that there is a widespread intuition that water is valuable, it is doubtful that there is a widespread intuition that water is essentially, necessarily, or finally valuable. Indeed, most people probably think that water is only contingently and instrumentally valuable.

The facts just reviewed create an unwelcome choice point for someone who wants to make an argument about water analogous to one of the more ambitious versions of Zagzebski's argument from intuition. She can start by invoking an intuition about water that a non-philosophical account could in principle vindicate – but then, in order to reach her desired conclusion that a philosophical theory of water should vindicate this intuition, she must insist that a philosophical theory of water that vindicates that intuition is better, all else equal, than one that allows the task to be accomplished elsewhere; and it is unclear how to justify this territorial attitude. Alternatively, she can start by invoking an intuition about water that philosophy is uniquely positioned to vindicate – but then, she makes a very dubious claim about our intuitions.

With this cautionary tale in mind, let's return to knowledge. The value intuition is an intuition to the effect that knowledge is more valuable than true belief. There are many ways that this could turn out to be true, each corresponding to a different way of explaining why the intuition is true. You could explain why knowledge is more valuable than true belief by giving a philosophical theory of knowledge that on its own entails that

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

knowledge is more valuable than true belief. You could also provide a theory of knowledge that, in conjunction with certain facts that do not figure in the theory itself – e.g., facts about our cognitive limitations, our practices of criticizing certain belief-forming methods, or our practical interests – explains why knowledge is more valuable than true belief. Finally, you could explain why knowledge is better than true belief purely in terms of facts that do not figure in your account of what knowledge is.

Given that all three styles of explanation are available, it is unclear why we should accept the central premise of the ambitious version of Zagzebski's argument from intuition: namely, that a philosophical theory of knowledge that explains (either in whole or in part) why knowledge is better than true belief is to be preferred, other things equal, to one that doesn't. For until we have some special reason to think that this explanatory task couldn't be accomplished elsewhere – a reason that the argument from intuition certainly doesn't provide – showing special favor to philosophical theories that accomplish the task simply amounts to an unmotivated reluctance to delegate.

The problem just outlined arises because the value intuition, as it is classically articulated, is one that philosophy is not uniquely positioned to vindicate. So, we might be tempted to try to fix the ambitious version of Zagzebski's argument by invoking a version of the value intuition that philosophy *is* uniquely positioned to vindicate: the intuition that knowledge is *essentially*, *necessarily*, or *finally* more valuable than true belief. But it is doubtful that there really is a widespread intuition with this content, just as it is doubtful that there is a widespread intuition that water is essentially, necessarily, or finally valuable.

I don't say this because I doubt that people have any intuitions about essences, necessity, or final value; it seems to me that we do have such intuitions, even if we wouldn't always use those terms to describe them. For instance, I think people have the intuition that bachelors are necessarily unmarried. The clearest symptom of this intuition is our response towards statements like, "It just so happens that bachelors are unmarried, but it could have been that bachelors were married." Such counterfactuals are borderline unintelligible; they elicit confusion, and the feeling that their speakers must not completely understand the terms that they are using. By contrast, corresponding counterfactuals about the value of knowledge – e.g., "If we were all omniscient, knowledge would be no more valuable than true belief" – do not strike us as gibberish. This suggests that there is not a widespread, pre-theoretical intuition to the effect that knowledge is essentially or necessarily more valuable than true belief.[3] (Note: This is not to say that the counterfactual test suggests that it is false that knowledge essentially or necessarily more valuable than true belief; it just suggests that if this is true, it is not a truth that we simply intuit.)

Likewise, I do think we have some intuitions about final value – or, to put it a bit more naturally, about things mattering in themselves. It's plausible, for instance, that there is a widespread intuition to the effect that happiness is good in itself, and not (just)

[3]Baehr (2009) and Botting (2020) also defend this claim, but in a way that fails to appreciate the genericity of the value intuition. They identify instances in which a bit of knowledge does not seem more valuable than the corresponding merely true belief, and they take these to show that people do not have the value intuition. But this line of argument would only work if generics embedded under a *essentially/necessarily* operator were truth-conditionally equivalent to universal generalizations embedded under those operators – and it is far from obvious that this is true. In other words, it is not obvious that *Essentially/necessarily, Fs are Gs* is truth-conditionally equivalent to *Essentially/necessarily, all Fs are Gs*. For a defense of this point and convincing illustrations of its importance for philosophical methodology, see Johnston (2016) and Johnston and Leslie (2012).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

because it conduces to some further end. But is it a widespread intuition that knowledge is not only valuable in itself but also *more* valuable in itself than true belief is valuable in itself, or perhaps valuable in itself while true belief is only valuable instrumentally? This strikes me as highly doubtful. I certainly do not have this intuition, nor do the non-philosophers whom I have asked. Of course, this is not to say that I am convinced that knowledge is *not* finally valuable, or in some sense more finally valuable than true belief. I am open to being convinced of these things, but I do not already intuit them.

If all that is right, then demands for ambitious vindicatory accommodation of the value intuition face the same problems as the corresponding demands for theories of water. As things stands, they invoke an intuition that philosophy is not uniquely positioned to vindicate, and so their key metatheoretical premises encode an odd, unmotivated territoriality. The problem could be fixed if the content of the motivating intuition were embedded under a *necessarily* or *essentially* operator, or if we articulated it in terms of final value – but then the resulting desideratum would rely on a very dubious claim about our intuitions.

1.2.2. More modest vindicatory accommodation

Demands for ambitious vindicatory accommodation of the value intuition run into trouble because the metatheoretical premise that animates them – much like the one animating demands for ambitious error-allowed accommodation of the value intuition – involves a distorted sense of which explanations are the special province of philosophy. But what about more modest versions of the demand for vindicatory accommodation of the value intuition? What if we just ask for a theory of knowledge that is compatible with a good explanation of why knowledge is more valuable than true belief, or with the fact that knowledge is more valuable than true belief?

I will argue that this reach for modesty topples us over into meaninglessness – which is just what happened, you'll recall, when we considered modest versions of the demand for error-allowed accommodation of the value intuition. So, desiderata centering vindicatory accommodation of the value intuition confront a familiar dilemma: either they rely on a dubious metatheoretical premise or they fail to generate meaningful constraints on theories of knowledge.

This second horn gets its point from the generic character of the value intuition. Naively, we might have thought that for a theory of knowledge to be compatible with the content of this intuition, it must not entail that there could be an instance of knowledge equally (or less) valuable than the corresponding true belief. But this is not so, because the value intuition is not a universal generalization; it is a generic. Just as *red wine is better than white* can be true even if some white wines are better than some reds, *knowledge is more valuable than true belief* can be true even if some instances of knowledge are not more valuable than the corresponding merely true beliefs.

Generics are truth-conditionally distinct not only from the corresponding *all* statements but also from the corresponding *most*, *usually*, and *normally* statements. Consider:

(1) Signed manuscripts are more valuable than unsigned ones.

There is a readily available interpretation of (1) that strikes us as true, even though it is not the case that most/normal signed manuscripts are more valuable than most/normal unsigned ones, or that most/normal manuscripts are such that a signature from the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

author confers added value. So, it would be a mistake to assume that (1) is incompatible with the claim that for most/normal manuscripts, a signature from the author does not confer added value. Likewise, it would be a mistake to assume that *knowledge is more valuable than true belief* is incompatible with the claim that most/normal instances of knowledge are no more valuable than the corresponding merely true beliefs.

Finally, certain generics seem to involve kind-level predication, and hence are compatible even with the falsity of the corresponding existential claim. Consider (2) as uttered in a conversation about whether to direct limited conservational resources towards sea otters or pandas:

(2) The sea otter is more valuable than the panda.

The intended interpretation of (2) is made true by the fact that sea otters, in contrast to pandas, are a keystone species of the ecosystem in which they participate. In other words, if sea otters were to disappear, this would have devastating consequences for the whole ecosystem that they inhabit, but the same is not true of pandas. Because the value associated with keystone-hood inheres in species rather than individual animals, this interpretation of (2) is true even if no individual sea otter is more valuable than some individual panda. The lesson here is that *knowledge is more valuable than true belief* might be true even if *no* individual bit of knowledge is more valuable than some true belief – because the relevant value might inhere in the kind, rather than individual instances of the kind.[4]

Besides illustrating a point about the flexibility of generics, this last example illustrates that the concept of *value* has a flexibility of its own. Something can be valuable instrumentally or for its own sake, contingently or necessarily, on its own or as part of a larger whole, in one respect or another; and all of these possibilities can be described simply by saying that the thing in question is valuable.

The upshot is that, thanks to the flexibility of generics and the concept of value, it is actually very difficult for a theory of knowledge to imply that the value intuition is mistaken. To pull this off, a theory of knowledge would have to entail that for every possible kind of value and for every possible interpretation of the generic, *knowledge is more valuable than true belief* is false. And I cannot think of a single theory of knowledge with this feature. What this suggests is that mere compatibility with the value intuition is a desideratum that is very easy to meet – so easy, in fact, that it is completely unhelpful as a guide to epistemological theorizing. Notice that this reasoning applies equally well to a slightly more ambitious version of vindicatory accommodation: compatibility with a good explanation of why the value intuition is true. Thanks to the flexibility of generics and the concept of value, there are not only many ways for the value intuition to be true but also just as many ways (and possibly more) to give a plausible explanation of why the value intuition is true. So, it is easy to achieve mere compatibility with a plausible explanation of why the value intuition is true.

In a way, I think, these lessons are hiding in plain sight in the literature on the value problem. Consider Zagzebski's influential critique of reliabilism. In brief, her argument goes like this: A good cup of coffee is not better for having been produced by a machine that reliably makes good coffee. Likewise, a true belief that has been produced by a process that reliably produces true beliefs is not better than a true belief that has happened by accident. According to reliabilism, knowledge is just true belief that has

[4]For an overview of the points about generics just reviewed, see Leslie and Lerner (2022).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

been produced by a process that reliably produces true beliefs. This means that, according to reliabilism, knowledge is not better than true belief. So, reliabilism contradicts a widespread intuition about knowledge, which is a serious count against it (Zagzebski 2003: 12–14).

Many have objected to this argument by pushing back on the idea that a reliably produced true belief is no better than its accidentally true counterpart. They point out that reliably produced true beliefs, like reliably produced cups of coffee, may be more instrumentally valuable with respect to certain ends than their flukey counterparts.[5] (As a coffee shop owner, you can charge more money for a coffee that is produced by a reliable machine, because the customer can be confident that the coffee will be good. That is at least one way in which reliably produced good coffee is better than accidentally good coffee.) Or they argue that even certain kinds of final value seem to depend on an object's history.[6] (Think of art objects, historical artifacts, and relics.) Or they note that one way of being valuable is meriting a positive evaluative attitude like admiration, and that whether a given true belief merits such an attitude might depend on how it was produced.[7]

All of these objections to Zagzebski make use of what I have called the *flexibility* of the concept of value. They show us that there are myriad ways for one thing to be more valuable than another, and that Zagzebski's argument against reliabilism ignores many of them.

We can also object to Zagzebski's argument by appealing to the generic character of the value intuition. Suppose that knowledge is reliably produced true belief, and that there is no way in which a reliably produced true belief is better than its accidentally true counterpart. Still, it might be true that knowledge is better than true belief – and therefore, we must acknowledge that reliabilism is compatible with the value intuition – since knowledge might be better than true belief in virtue of facts about the kinds rather than facts about particular instances of knowledge and true belief.

Taken individually, each of the above arguments offers a reason to think that reliabilism is compatible with the value intuition after all. Under this description, they threaten Zagzebski's argument against reliabilism, but they don't directly undermine the metatheoretical premise that animates that argument: namely, that compatibility with the value intuition is a substantive desideratum that provides a helpful starting point for building, criticizing, and refining theories of knowledge. Taken together, however, the various objections to Zagzebski's argument threaten that metatheoretical premise, too. Why? Because they provide a compelling illustration of the hypothesis that I've been advancing in this section: thanks to the genericity of the value intuition and the capaciousness of the concept of value, it is extremely difficult to see how *any* minimally plausible theory of knowledge could be incompatible with the value intuition, or rule out all plausible explanations of why the value intuition is true. And if that's right, then compatibility with the value intuition (or some plausible explanation of why that intuition is true) is so easy to achieve as to be completely unhelpful as a guide to epistemological theorizing.

Here again, we may be tempted to try to fix this problem by fiddling with the value intuition. Perhaps, we might suggest, its content is not a multiply ambiguous generic value claim that no theory of knowledge on the market contradicts; perhaps its content is something more tractable like *necessarily, every merely true belief would be made more valuable (in some particular respect) by being converted into knowledge*. If that were true,

[5]See, for instance, Olsson (2007).

[6]See, for instance, Brogaard (2006).

[7]See, for instance, Brady (2006), Riggs (2002), and Greco (2003).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

then compatibility with the value intuition would not be an inert desideratum, for it is at least conceivable that a minimally attractive theory of knowledge might contradict its precise, unambiguous content. But there is just no reason to think that most people – or even most philosophers – have an intuition with this content.

2. The argument from interest

For the reasons given above, I doubt that an appeal to intuition can motivate treating the value problem as a starting point for epistemological theorizing. But arguments from intuition are not the only way that philosophers have defended the metaepistemological significance of the value problem. There is also a family of arguments that we can call *arguments from interest*:

> The question *What is knowledge?* is not independent of the question *Why do we value knowledge?* For those who consider the former question prior, compare the pair of questions *What is knowledge?* and *Can we get it?* It is common for anti-skeptic naturalist epistemologists to say that whatever knowledge is, it has to be defined as something we have. We are not interested in a non-existent phenomenon. I say that knowledge has to be defined as something we value. We are not interested in a phenomenon with little or no value. (Zagzebski 2003: 26)

> One of the questions that is very rarely asked in epistemology concerns what is perhaps the most central issue for this area of philosophy. It is this: Why should we care about whether or not we have knowledge? Put another way: Is knowledge valuable and, if so, why? The importance of this question resides in the fact that it could well be that it is only if the primary focus of epistemological theorizing – i.e., knowledge – is valuable that the epistemological enterprise is itself a worthwhile undertaking. (Pritchard 2007: 37)

> [I]f knowlege is not the overarching epistemic telos with regard to particular propositions, why [has] such tremendous emphasis . . . been placed on the theory of knowledge in the history of philsophy, and just what function [does] that notion serve within that history? If knowledge is not the overarching purpose of inquiry, then why is the notion important, and why should we continue to be concerned in normative epistemology above all with what knowledge is and how it can be achieved? (Sartwell 1992: 174–5)

The argument expressed in these passages is somewhat difficult to pin down. For my purposes, however, we do not need a perfectly perspicuous articulation of the whole argument, because I want to focus on just one of its premises. This premise says: If knowledge isn't valuable, then the question *What is knowledge?* is (or "could well be," if we prefer Pritchard's hedge) uninteresting, i.e., unworthy of philosophical investigation. I'll call this the *interest premise*.

2.1. The interest premise requires defense

Strikingly, appeals to the interest premise generally do not include explicit defenses of this premise. In fact, I have not been able to find any defenses in published work –

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

although one has been suggested to me several times informally, and I will come on to this defense shortly. First, though, I would like to explain why the interest premise requires defense.

To see why we need a special reason to accept the interest premise, we can start by noticing that generalizing it yields an extremely surprising claim that most of us would probably reject. This is the claim that for all *X*, if *X* isn't valuable, the question *What is X?* is unworthy of philosophical investigation. Think about what this claim implies for, say, central research projects in metaphysics. Many metaphysicians are interested in the question *What is causation?* Do we really want to say that their investigations are only worthwhile if causation – the phenomenon itself, not the corresponding concept – is valuable? I think not. I am not even sure that it makes sense to say that causation is valuable; the claim seems to involve a category mistake. Nevertheless, I am confident that my colleagues working on causation are engaged in a worthwhile project. Or, think about possibility and necessity. I am confident that these things are worthy of philosophical study. But I don't think that possibility and necessity are valuable – and here again I am talking about the relations themselves, rather than the concepts that pick them out.

These examples tell against the idea that in general, a phenomenon is only worthy of philosophical analysis if it is valuable. So, a defender of the interest premise shoulders the explanatory burden of showing that knowledge in particular, in contrast to things like causation and modality, is such that it must be valuable in order to deserve philosophical study.

Before moving on to consider how this might be done, I'd like to say something about an alternative to the interest premise that the above reflections on causation and modality naturally bring to mind. This is the claim that if the *concept* of knowledge isn't valuable, then the question *What is knowledge?* is unworthy of philosophical investigation. Call this the *conceptual interest premise*. On its face, the conceptual interest premise is much more plausible than the interest premise, because the more general idea of which it is an application – for all *X*, if the concept of *X* isn't valuable, then the question *What is X?* is uninteresting – is much more plausible than the generalization of the original interest premise. We can see this by reflecting on the cases that cause trouble for that latter generalization: even if causation and modality aren't themselves valuable, the concepts of *cause*, *possible*, and *necessary* surely are.

But epistemologists driven by the value problem can take no comfort from these observations. Why not? Because if we construct a version of the argument from interest using the conceptual interest premise in place of the interest premise, we arrive at the conclusion that the concept of knowledge should be defined as something valuable; we don't arrive at Zagzebski's claim that knowledge itself must be defined as something valuable. And there is a big difference between these two claims since, as we have just seen, the concept of *X* may be valuable even if *X* itself is not. So, an account of knowledge that assumes that the concept of knowledge is valuable could succeed by its own lights without implying that knowledge is valuable; in fact, it could succeed while implying that knowledge is not valuable! It would therefore be a mistake to think of the conceptual interest premise as a variant of the interest premise or as necessarily connected to that premise in any way. Relatedly, it would be a mistake to understand epistemology driven by the conceptual interest premise as offering an attractive version of theorizing driven by the value problem, or as vindicating any of the metatheoretical premises that drive such theorizing.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

These points may be obvious, but they bear stressing because there is a substantial research program often discussed in connection with the value problem that arguably treats the conceptual interest premise as the core tenet of its methodology. We can call it the *function-first* approach. Epistemologists working in this tradition begin their inquiry into what knowledge is by asking what valuable purpose the concept of knowledge serves for us. Then they essentially "reverse-engineer" the definition of the concept to ensure that the concept is fit for purpose. Craig (1999) offers a classic contribution to this research program; more recent function-firsters include Hannon (2014), Kelp (2011), and Beebe (2012).

Some function-first epistemologists are not only interested in explaining why the concept of knowledge is valuable; they are also interested in shedding light on the apparent value of knowledge itself. This is why function-first epistemology gets discussed in connection with the value problem,[8] and why it might be briefly tempting to think that function-first epistemology *itself* is driven by the value problem, or that the relatively plausible metatheoretical premises underpinning the function-first approach are somehow connected to the less plausible ones endorsed by value-driven theorists like Zagzebski. But as I have argued, these are mistakes. The concept of knowledge could serve a valuable function even if knowledge itself were worthless; relatedly, a theory could explain why the concept of knowledge is useful even while implying that knowledge itself is useless. We therefore cannot treat function-first epistemology as containing or vindicating anything meaningfully similar to the argument from interest, nor can we think of the conceptual interest premise as a close relative of the interest premise.

2.2. Defending the interest premise by appeal to "normativity"

Back to the interest premise itself, then. If the interest premise requires defense, as I have argued, how might we defend it?

Doubtless there are many strategies worth exploring here. But I will leave the bulk of this exploration to others; there is just one defense of the interest premise that I'd like to discuss before concluding. I've chosen this argument because a number of people have sketched something like it when I've asked how they might defend the interest premise – though I have not seen it laid out explicitly in print anywhere. Here is how one anonymous respondent put the argument:

> The idea that we must vindicate the value intuition could come from the concern that if knowledge isn't shown to be valuable then why should anyone care about the type of normativity around structured around it? Might then as well inquire into the norms of etiquette or some other arbitrary system of norms not grounded in any real value. But epistemology is supposed to be more important than that type of enterprise.

Let's call this the *normativity argument*. It begins with the observation that there is a system of norms, rules, and/or evaluative standards that orbit around knowledge.

[8]See, for instance, §6 of Pritchard *et al.* (2022).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

For instance, it has been proposed that you should believe p only if that is justified by what you know;[9] assert p only if you know p;[10] act as if p only if you know p;[11] and use p as a premise in practical reasoning only if you know p.[12] It could also be that epistemic rationality, justification, and/or reasonableness are best understood in terms of their relationship with knowledge. Let's sum up these points by saying that there is a normativity associated with knowledge.

The second step of the argument is drawing a distinction between so-called *robust* and merely *formal* normativity. What exactly this contrast comes to is a more difficult matter than is sometimes acknowledged, I think. But we can get a rough handle on it, and one that will do for our purposes, by reflecting on claims like: "I get that I should ϕ according to the norms of X, but I have no real reason to follow the norms of X and so no real reason to ϕ." If you can imagine a realistic situation in which it would make sense to say this about a certain set of norms – e.g., the norms of etiquette – then that set of norms is paradigmatically formal. If you struggle to imagine such a situation, as you might if we are discussing, say, the norms of morality, then that is a sign that the norms under discussion are robust.

The third premise of the argument is a claim about what is and isn't worthy of philosophical study: Insofar as philosophical inquiries into purely formal norms are worth pursuing at all, they are less important than inquiries into robust norms. And the fourth and final premise is that if knowledge isn't valuable, then the normativity associated with it must be purely formal normativity. Putting all this together, we get a slightly more modest but still quite substantial version of the interest premise: If knowledge isn't valuable, then philosophical investigation into the norms associated with it is less important than we might have thought or hoped.

There are a number of places where we might put pressure on this argument, but I'd like to focus on a kind of dilemma that arises for anyone who wants to endorse the third and fourth premises together. We can set up this dilemma by posing the question: Is the norm that says "Don't feel guilty unless you've done wrong" robust or purely formal?

Suppose we say that it is robust. Then we can construct a "companions in innocence" argument against the fourth premise, along the following lines: If the norm *Don't feel guilty unless you've done wrong* is underpinned by some kind of value, it is certainly not the value of having done wrong. That has no value, presumably. The general lesson here is that the fittingness conditions for various attitudes and actions are sometimes spelled out in terms of things that are not valuable, which (we are supposing on this horn of the dilemma) does not undermine the robustness of the corresponding norms. So, even if knowledge is not valuable, norms like *Don't assert p unless you know that p* – which, you'll notice, is structurally just like the norm *Don't feel guilty unless you've done wrong* – might still be perfectly robust.

Now suppose that the latter norm is purely formal. On this horn of the dilemma, we draw the formal/robust distinction in such a way that norms describing when guilt is fitting fall on the formal side of the line. But if we do this, we seriously undermine the claim that formal norms are unworthy of philosophical study – the third premise of the normativity argument. Why? Not only because that claim now awkwardly implies that many of our colleagues working in moral psychology are wasting their time, but also and

[9]Williamson (2000: ch. 9).
[10]Williamson (2000: ch. 11).
[11]Fantl and McGrath (2002).
[12]Hawthorne and Stanley (2008).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

more importantly, because the many reasons that we have to care about when guilt is fitting (and the philosophical challenges that we face in trying to understand those reasons) do not go away simply because we have deemed the corresponding norms "purely formal." These reasons and challenges are not spelled out in terms of the formal/robust distinction, and so they do not depend on how exactly we draw that distinction.

We can also attack the third premise of the normativity argument more directly. Notice that even if we don't always have reason to care about being in accord with formal norms, certainly we sometimes do. For example, if you've promised your niece to play chess with her, you have reason to comply with the formal norms of chess. This alone doesn't mean that the norms of chess are worthy of philosophical study. They aren't because (a) we know what the rules of chess are and (b) there is no special mystery about when we have reason to comply with them or why we have reason to comply with them when we do. But sometimes matters are not so clear. So, unless we decide to draw the formal/robust distinction in such a way that formal norms are simply defined as philosophically perspicuous – an odd choice, surely – we must allow that even if a species of normativity is purely formal, there may still be deep puzzles surrounding what the relevant norms are, when we do and don't have genuine reason to comply with them, and why we have reason to comply with them when we do. In other words, we must grant that a norm's being formal does not mean that it is unimportant, or that it raises no distinctively philosophical puzzles – which is to say, we must give up the idea that formal norms as such are unworthy of philosophical investigation.

It is particularly important to keep these points in mind when we are discussing epistemic norms, because there are some epistemologists who argue that these norms are purely formal while also demonstrating their commitment to the idea that these norms are worthy of philosophical study. This is a wholly coherent position. Why? Because even if epistemic norms are purely formal, as these thinkers define the notion, that does not dissolve any of the deep puzzles about what the norms are, when we do and don't have genuine reason to comply with them, and why we have reason to comply with them when we do. It merely places certain constraints on how those puzzles may be solved.[13]

For the reasons just reviewed, I don't think that the normativity argument can work. Perhaps we can find some other way to defend some version of the interest premise – but I will leave the search to those who are more optimistic about it than I am.

3. Conclusion

There may be a widespread intuition to the effect that knowledge is more valuable than true belief, and this intuition may be correct. But even if these things are true, they provide little threat to existing theories of knowledge and little help to those who would like to build better alternatives.

References

Baehr, J. (2009). 'Is There a Value Problem?' In A. Haddock, A. Millar and D. Pritchard (eds), *Epistemic Value*, pp. 42–59. Oxford: Oxford University Press.

Beebe, J. (2012). 'Social Functions of Knowledge Attributions.' In J. Brown and M. Gerken (eds), *Knowledge Ascriptions*, pp. 220–42. Oxford: Oxford University Press.

[13]See Maguire and Woods (2020) for an example of a view according to which epistemic norms are purely formal in some sense but still worth caring about.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137

Botting, D. (2020). 'The Value Problem of A Priori Knowledge.' *Acta Analytica* **35**, 229–52.
Brady, M. (2006). 'Appropriate Attitudes and the Value Problem.' *American Philosophical Quarterly* **43**(1), 91–9.
Brogaard, B. (2006). 'Can Virtue Reliabilism Explain the Value of Knowledge?' *Canadian Journal of Philosophy* **36**(3), 335–54.
Craig, E. (1999). *Knowledge and the State of Nature*. Oxford: Oxford University Press.
Fantl, J. and McGrath, M. (2002). 'Evidence, Pragmatics, and Justification.' *The Philosophical Review* **111**(1), 67–94.
Greco, J. (2003). 'Knowledge as Credit for True Belief.' In M. DePaul and L. Zagzebski (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, pp. 111–34. Oxford: Oxford University Press.
Hannon, M. (2014). 'Fallibilism and the Value of Knowledge.' *Synthese* **191**(6), 1119–46.
Hawthorne, J. and Stanley, J. (2008). 'Knowledge and Action.' *Journal of Philosophy* **105**(10), 571–90.
Johnston, M. (2016). 'Remnant Persons: Animalism's Undoing.' In S. Blatti and P.F. Snowden (eds), *Animalism: New Essays on Persons, Animals, and Identity*, pp. 89–127. Oxford: Oxford Academic.
Johnston, M. and Leslie, S. (2012). 'Concepts, Analysis, Generics, and the Canberra Plan.' *Philosophical Perspectives* **26**, 113–71.
Jones, W. (1997). 'Why do we Value Knowledge?' *American Philosophical Quarterly* **34**(4), 423–40.
Kelp, C. (2011). 'What's the Point of 'Knowledge' Anyway?' *Episteme* **8**(1), 53–66.
Leslie, S. and Lerner, A. (2022). 'Generic Generalizations.' In E.N. Zalta and U. Nodelman (eds), *The Stanford Encyclopedia of Philosophy*. Stanford: The Metaphysics Research Lab.
Maguire, B. and Woods, J. (2020). 'The Game of Belief.' *Philosophical Review* **129**(2), 211–49.
Meylan, A. (2007). The Value Problem of Knowledge: [A]gainst a Reliabilist Solution. *Proceedings of the Latin Meeting in Analytic Philosophy Genova*, 85–92.
Olsson, E. (2007). 'Reliability, Stability, and the Value of Knowledge.' *American Philosophical Quarterly* **44**(4), 343–55.
Pritchard, D. (2009). 'Knowledge, Understanding, and Epistemic Value.' *Royal Institute of Philosophy Supplement* **64**, 19–43.
Pritchard, D. (2007). 'The Value of Knowledge.' *Aporia* **1**, 37–45.
Pritchard, D., Turri, J. and Carter, J.A. (2022). 'The Value of Knowledge.' In E.N. Zalta and U. Nodelman (eds), *The Stanford Encyclopedia of Philosophy*. Stanford: The Metaphysics Research Lab.
Riggs, W. (2002). 'Reliability and the Value of Knowledge.' *Philosophy and Phenomenological Research* **64**(1), 79–96.
Sartwell, C. (1992). 'Why Knowledge is Merely True Belief.' *Journal of Philosophy* **89**, 167–80.
Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
Zagzebski, L. (2003). 'The Search for the Source of Epistemic Good.' *Metaphilosophy* **34**(1), 12–28.

Cite this article: Dandelet S. (2026). "Metaepistemology and the Value Problem." *Episteme* 1–15. https://doi.org/10.1017/epi.2026.10137

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 11:16:55, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/epi.2026.10137