

Safety's Swamp: Against the Value of Modal Stability

GEORGI GARDINER

Rutgers University

Final version appears in *American Philosophical Quarterly*

Abstract

An account of the nature of knowledge must explain the value of knowledge. I argue that modal conditions, such as safety and sensitivity, do not confer value on a belief and so any account of knowledge that posits a modal condition as a fundamental constituent cannot vindicate widely held claims about the value of knowledge. I explain the implications of this for epistemology: We must either eschew modal conditions as a fundamental constituent of knowledge, or retain the modal conditions but concede that knowledge is not more valuable than that which falls short of knowledge. This second horn—concluding that knowledge has no distinctive value—is unappealing since it renders puzzling why so much epistemological theorising focuses on knowledge, and why knowledge seems so important.

Introduction

Epistemology sometimes aims to understand the nature of knowledge by finding the individually necessary and jointly sufficient conditions which capture our judgements about the extension of knowledge. That is, they hope to offer a set of conditions such that all and only those cases that satisfy the conditions are instances of knowledge. In Edward Craig's terms, they seek 'an explicit intension to match the intuitive extension'.¹ But matching the extension is not enough: In addition to picking out the extension of knowledge, these conditions should explain the nature, role, and value of knowledge. They should explain, in other words, what knowledge is, what it does, and why it is significant. Claims about the value of knowledge are thus a constraint on theories about the nature of knowledge.

In this paper I articulate widely held claims about the value of knowledge, and argue that any account that posits modal conditions as a fundamental constituent on knowledge cannot vindicate these value claims. This is because, I argue, the modal condition's obtaining does not confer any value on a belief; the value of the modal condition is swamped by the value of other features of the belief. I articulate the implications of this for theorising about knowledge: We must either eschew modal conditions on knowledge, or retain the modal conditions but concede that knowledge is not more valuable than that which falls short of knowledge. This second horn—concluding that knowledge has no distinctive value—is unappealing since it renders puzzling why so much epistemological theorising focuses on knowledge, and why knowledge seems so important to us.

¹ Craig (1990) p. 1.

Two relatively uncontroversial necessary conditions on knowledge are truth and belief; it is widely held that S knows that p only if p is true and S believes p.² These conditions help to pick out the extension of knowledge and help explain its nature, role, and value. But truth and belief are not jointly sufficient to pick out the extension—many true beliefs are not known. The results of lucky guesses, for example, are not knowledge. The conjunction of truth and belief, furthermore, cannot adequately explain all of the functions and value of knowledge.³ To further illuminate the contours of knowledge, epistemologists aim to formulate conditions that, when combined with true belief, constitute knowledge. Conditions that have been proposed include that the belief is justified, that the belief is formed through cognitive virtue, that the person stands in the right causal relationship to the truth of the belief, that the person can rule out all relevant alternatives, that the belief was reliably formed, that the belief is sensitive, and that the belief is safe. These last two conditions, safety and sensitivity, are known as modal conditions because they concern counterfactual properties of the belief and believer.

Recent epistemological theorising features many intricate accounts of knowledge that attempt to capture intuitive knowledge attributions in a range of hard cases. In response to the plethora of (increasingly complex and convoluted) accounts proposed in the literature, Jonathan Kvanvig emphasises that any account of the nature of knowledge must accord with claims about the value of knowledge.⁴ Claims about the value of knowledge, in other words, restrict which conditions on the nature of knowledge are plausible.

In this paper I argue that widely accepted claims about the value of knowledge are in tension with the view that an analysis of knowledge must include a modal condition, such as safety or sensitivity, as a fundamental constituent. This is because modal conditions do not add value to a true belief.

The Value Claims

We seem to care a great deal about knowledge. As society and individuals, we invest substantial efforts gaining and disseminating knowledge. We seem to value knowledge for its own sake, and knowledge is important for practical matters; we tend to act on and assert only what we know. The term “knows” and its cognates appear often in conversation and “knows” is one of the few terms that exists in every language.⁵ In philosophy and other disciplines the nature of knowledge has long been an important research question. This all suggests that knowledge has some special value.

The primary value claim, which arguably dates back to Plato’s *Meno*, holds that knowledge is more valuable than mere true belief. This is a relatively weak claim—it holds simply that it is (normally) better to know p than to merely believe p correctly. The secondary value claim, which is strictly stronger, holds that knowledge is more valuable than any proper subset of its parts. If knowledge is a whole comprising {X, Y, Z}, then the

² Nothing in philosophy is wholly uncontroversial: Hazlett (2010) challenges the view that knowledge is factive. Lewis (1996) and Radford (1966) challenge the view that knowledge entails belief.

³ Suppose someone tells you p, for example, and you rely on his testimony. You would be vexed if you later learnt he simply guessed that p, even if you also learn that p is true.

⁴ Kvanvig (2003) p.6. Independent recent statements of the idea that an account of knowledge must offer resources to explain the value of knowledge can be found in Zagzebski (1999), Sosa (2000), and Williamson (2000), chapter one.

⁵ The term “knows” is the eighth most commonly used verb in English, according to the *Oxford English Dictionary*. See also Goddard (2010).

secondary value claim holds that any proper subset, such as {X, Y}, has less value than {X, Y, Z}. If knowledge were ‘justified, safe, true belief’, for example, ‘justified true belief’ would be a proper subset of its parts, and—according to the secondary value claim—would be less valuable.⁶

These value claims are weak, and are correspondingly plausible and widely held. If these value claims are false it is hard to explain why we seem to care so much about knowledge, why people value attaining and sharing knowledge, why knowledge rather than other cognitive states plays such a large role in our lives, and why so much epistemological focus has been concerned with knowledge, rather than competing epistemic states such as true belief. The two value claims can be seen as constraints on a theory of knowledge: To the extent that the claims are plausible, any proposed theory must either explain, or explain away, these value claims.

In this paper I argue that those who posit safety or sensitivity as a necessary condition on knowledge are severely restricted in what value claims they can make about knowledge. This is because, I argue, being safe (or sensitive) does not confer any additional value onto a true belief. I focus on safety, since it is the more favoured modal condition, but my arguments apply to all merely modal conditions. My thesis thus generalises: the obtaining of modal conditions does not confer value onto a belief.

In what follows I describe the safety condition and how it can feature in an account of knowledge. I then examine four attempts to establish safety’s value, and I explain why safety cannot yield the value these authors suggest. Finally, I summarise the implications for modal epistemology.

Safety

Safety and sensitivity conditions on knowledge attempt to explicate the view that knowledge is immune from luck—S’s knowing that p is incompatible with S’s belief being true by accident—and there must be some non-chance connection between the belief that p and the fact that p. Robert Nozick advanced the sensitivity condition on knowledge. The sensitivity condition holds that S knows that p only if S’s belief that p is sensitive. S’s belief that p is sensitive iff if it were not the case that p, S would not believe that p.⁷ The sensitivity condition on knowledge has been subject to various criticisms and currently garners relatively few proponents.⁸

An alternative modal condition, the safety condition, was first advanced separately by Ernest Sosa and Timothy Williamson, and has been developed Duncan Pritchard.⁹ The safety condition holds that S knows that p only if S’s belief that p is safe. S’s belief is safe iff were S to believe that p, p would not be false. In other words S’s belief that p is safe iff S’s belief that p could not easily have been wrong. When introducing the safety condition Sosa writes,

⁶ For discussion of these value claims, see Kvanvig (2003), Greco (2009), Riggs (2009), Pritchard (2007; 2009a), and Pritchard, Millar and Haddock (2010). The claims are also known as Meno assumptions, after Plato’s dialogue in which the value of “episteme”—often translated as “knowledge”—is discussed.

⁷ Nozick (1981). Given a Lewisian understanding of counterfactual conditionals, sensitivity is equivalent to the requirement that in the nearest possible worlds in which p is false, the subject does not believe p.

⁸ For discussion, see Sosa (2000); Vogel (1987); Kripke (2011). See also DeRose (1995; 2010); Roush (2005).

⁹ Sosa (1999); Williamson (2000), esp. pp. 159-60; Pritchard (2005). For recent criticisms of the view that knowledge entails safety, see Comesaña (2005) and Baumann (2008).

S's belief that *p* is "safe" iff: S would believe that *p* only if it were so that *p*. (Alternatively, a belief by S that *p* is "safe" iff: S would not believe that *p* without it being the case that *p*; or, better, iff: as a matter of fact, though perhaps not as a matter of strict necessity, not easily would S believe that *p* without it being the case that *p*.)¹⁰

Given a Lewisian understanding of counterfactual conditionals, safety is equivalent to the requirement that, in the nearest possible worlds in which S believes *p*, *p* is true.¹¹ This possible worlds interpretation of safety is made explicit by Pritchard:

S's belief is safe iff in most near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which S continues to form her belief about the target proposition in the same way as the actual world, her belief continues to be true.¹²

According to a Lewisian possible worlds interpretation of safety, to determine whether the person's belief that *p* is safe we must consider the nearest possible worlds where the person also believes *p*. If there are relevantly similar worlds where the person believes that *p*, yet her belief is false, then her belief is not safe in the actual world. Facts about the actual world, such as the epistemic friendliness of her environment, her epistemic abilities, and her evidence, determine which worlds are nearby. (Facts such as what the person had for breakfast or what flowers she planted in her garden will also determine which worlds are more similar, but these are unlikely to be epistemically relevant.) If a person's belief is true only by chance because, for example, her belief constitutes a Gettier case or a lucky guess, then typically in many of the relevantly similar worlds, her counterpart does not enjoy similar fortune, and the counterpart's belief is false; in the actual world, her belief is unsafe.¹³

The safety condition can feature in an account of knowledge in two ways. Robust anti-luck epistemology maintains that knowledge is safe true belief. No other conditions are required for an account of knowledge.¹⁴ Weak anti-luck epistemology maintains that safety is a necessary condition on knowledge, but contends that safe true belief is not sufficient for knowledge. According to weak anti-luck epistemology, knowledge is safe true belief plus other condition(s), such as that the belief is justified, rational, virtuously held, reliably formed, or that the agent has the right causal connection with the facts.¹⁵

¹⁰ Sosa (1999).

¹¹ Lewis (1973).

¹² Pritchard (2009b), p. 34. This is Pritchard's favoured version of safety. It exhibits a two-level requirement: The nearest worlds are weighted as more significant and so the belief must be true in all the nearest worlds. More false beliefs are permitted in worlds slightly more distant (Pritchard 2009a, p. 37). The version also makes explicit basis-relativity in the belief-forming method. These details need not concern us here, since my arguments apply to any variant of safety.

¹³ True beliefs about necessary propositions are true in nearby worlds, even if the belief is only true by chance.

¹⁴ *cf.* Pritchard (2005).

¹⁵ *cf.* Pritchard (2009a); Pritchard (2012).

Safety and Value

This discussion raises the question of what value, if any, safety contributes to a true belief. There are two ways that knowledge can have greater value than lesser epistemic standings. One is that the constituent parts each have some value that the part contributes to the whole. This mechanism is consistent with the additivity of value; the value of the sum of the parts is equal to the sum of the value of the parts. The other is that when the parts of the whole fuse in the right way, value emerges which is either more than the sum of the parts, or different in kind from the value of the parts.¹⁶ With either mechanism, if the value of knowledge is greater than the value of any proper subset of its parts—in other words, if the secondary value claim is correct—then removing a part from the whole would reduce the value of the remainder. The value of a part cannot be *swamped* by the value of another part. Swamping occurs when all the value contributed by a constituent is already present in the rest of the whole.¹⁷ A familiar putative example of swamping concerns the value of ‘being reliably formed’ in the whole ‘reliably formed true belief’. Many argue that ‘being reliably formed’ contributes no unswamped value to true belief. They argue that being reliably formed is only valuable insofar as we value discovering the truth, and so the value of the belief’s being true swamps the value of its being reliably formed.

The swamping mechanism can be generalised: If we value something, M, only instrumentally in virtue of something else, N, and N is present then in general the value of M is swamped by the presence of N. M confers no additional value onto the whole.

A modal condition on knowledge cannot vindicate the value claims articulated in the second section of this essay. This is because, I suggest, safety confers no unswamped value on true belief. This means that anyone who holds robust anti-luck epistemology, the view that knowledge is safe true belief, is committed to the falsity of the primary and secondary value claims. Robust anti-luck epistemology cannot explain why knowledge is more valuable than mere true belief, and so cannot explain why knowledge is more valuable than any subset of its parts. Anyone who holds weak anti-luck epistemology, the view that knowledge is safe true belief plus some other conditions, can vindicate the primary value claim, since the other conditions on knowledge might confer additional value on a true belief. They must, however, deny the secondary value claim, that knowledge is more valuable than a proper subset of its parts. This is because if safety adds no value to a true belief, as I argue, then a whole comprising a proper subset of the conditions of knowledge, but where that belief is not safe, has the same value as knowledge.

Below I explore four attempts to establish that the obtaining of modal conditions contributes value to a true belief, and explain why each is unsuccessful.

First Strategy: Enabling Further True Beliefs

The first attempt is advanced by Pritchard. Pritchard suggests that safety adds value to a true belief because,

After all, one might plausibly argue that [safety] will, *ceteris paribus*, better enable you to have further true beliefs in the future. Thus, relative to the epistemic good of true belief, it does seem epistemically valuable.¹⁸

¹⁶ Moore (1903).

¹⁷ See especially Kvanvig (2003), who credits Swinburne (1999; 2000).

¹⁸ Duncan Pritchard, personal correspondence, February 2009.

Pritchard argues that an individual belief's being safe will, *ceteris paribus*, better enable an agent to have further true beliefs. But this is implausible—nothing about a belief's safety as such bears on the truth of other beliefs. To see this, consider a thought like Descartes's *cogito*, which says, "I am thinking". It is safe, but its safety as such does not enable further true beliefs. Its infallibility performs all the dialectical roles. Or consider the following case: A demon changes the world to ensure Olivia's beliefs are maximally false, except for one belief, about which he changes the world to make sure it is true. This true belief is safe, since Olivia cannot believe it without it being true—and so in all the nearby worlds that Olivia believes that *p*, *p* is true—but the safety of this one belief has no effect on the amount of true beliefs that Olivia holds.¹⁹

Whilst this may seem like an exotic and hence marginal case, it illustrates the relevant characteristic of safety—safety concerns only the counterfactual properties of the belief in question. A belief's merely being safe does not affect whether the person's future beliefs will be true. This is a problem for Pritchard's proposal about the value of safety. I return to this point in my discussion of the third strategy for establishing the value of safety. One could respond that safety is valuable because it normally enables more true beliefs, even if there are some exceptions. Cases such as the Olivia case or the *cogito* are not a threat to this proposal because they are abnormal cases. But as I explain below and in my discussion of the third strategy, it is not the safety as such that normally enables further true beliefs. It is the properties that give rise to the safety.

An alternative response holds that a belief's being safe, rather than merely true, leads to further true beliefs in the future because i) a person is more likely to use the same belief-forming process, evidence-type, or cognitive abilities again if it does not lead merely to a true belief, but a *safe* true belief, or ii) a safe belief, as opposed to a merely true belief, is likely to be held with more conviction, and so 'stay put' in the face of potential defeaters. This understanding of Pritchard's suggestion evades the worry articulated above because, according to this suggestion, safety does influence other beliefs, albeit indirectly.

These two suggestions are analogous to reliabilist responses to the value problem posed by the swamping mechanism discussed above. The swamping problem for reliabilism is that the value of true belief swamps the value of being reliably formed and so reliabilism, which holds that knowledge is reliably formed true belief, cannot account for the value of knowledge over mere true belief. One reliabilist strategy for overcoming this problem argues that the value of a belief's being formed by a reliable method stems from the increased likelihood of future beliefs being true.²⁰

This interpretation of Pritchard's suggestion would not suit the modal epistemologist, however. Safety is an externalist condition—the condition's obtaining need not be accessible to the agent. The success of the two strategies above requires an internalist condition to

¹⁹ Similarly we can imagine a case where the demon changes the world to make sure all Olivia's beliefs are true, or half of them are true. In all of these cases Olivia's true beliefs are safe, since Olivia cannot believe *p* unless *p*, but the mere safety of her beliefs has no effect on the number of true beliefs Olivia holds.

²⁰ Goldman and Olsson (2009). Goldman and Olsson write, "if you have used a given method before and the result has been unobjectionable, you are likely to use it again on a similar occasion, if it is available." Goldman and Olsson refer to this as the "learning assumption". Similarly reliabilism can adopt the suggestion that a person might well have increased confidence in those beliefs she judges to be reliably formed, and so retain the belief in the face of defeaters.

obtain. The agent must detect that the belief is safe, and so alter her doxastic behaviour accordingly, by deploying the same belief-forming method again or retaining the belief in the face of potential defeaters. This internalist condition is not part of the safety condition, and so these strategies locate the value not with safety itself, but with some other condition. Since the extra condition is internalist rather than externalist, furthermore, it is in tension with the spirit of modal epistemology.

A modal epistemologist could respond that the relevant mechanism is not that the agent has second-order beliefs tracking the safety of her belief and adjusting her confidence or doxastic behaviour accordingly, but rather that an evolutionary mechanism is at play. According to this proposal, epistemic agents need not be aware of which beliefs are safe, but rather safe beliefs are selected for, and so lead to more true beliefs in the future. This suggestion qualifies as an externalist elucidation of how safe beliefs better enable further true beliefs. But this suggestion will not work: whether a belief is safe concerns non-actual worlds, and so safety as such cannot be selected for. Arguably reliability can be selected for.²¹ Reliability concerns the actual world over time, and so may be a candidate for an evolutionary hypothesis. But safety concerns nearby possible worlds and counterfactual situations. That a belief is true in nearby worlds, or would also have been true in counterfactual situations, cannot be selected for in the actual world. Thus it is false that safety is valuable because it enables more true beliefs in the future.

Second Strategy: The Relation to Other Valuable Properties

An alternative strategy for establishing the value of safety is noting that i.) if a belief is safe, it must have been well-formed, i.e. through virtuous epistemic agency, with an appropriate process, in a good environment, etc., and ii.) these things are valuable. The strategy holds that safe true beliefs are more valuable than mere true beliefs because if beliefs are safe, these other valuable properties are present.²²

One problem with this strategy is that we can imagine safe beliefs that do not have any of these valuable properties—in the Olivia case, for example, the agent does not stand in the right causal relationship with the truth of her belief, need not be a virtuous epistemic agent, need not believe according to the evidence etc. Her beliefs would still be safe. So safety does not always co-instantiate with other valuable properties. Thus if safety is valuable solely because it occurs in tandem with independently valuable properties, it is not ‘universally’ or always valuable.

Some epistemologists hold that the value constraints on a theory of knowledge concern every individual item of knowledge: The primary value claim should be understood as holding that for every proposition, *p*, it is better to know *p* than to merely truly believe *p*. The secondary value claim can be understood accordingly as holding that for every proposition, *p*, it is better to know *p* than to merely stand in some relation comprising elements of the knowledge relation.²³ To the extent this is a plausible understanding of the value claims, the fact that safety can diverge from these other bearers of value is a problem for this strategy.

But this is not the only difficulty with this strategy. Even when safety does co-instantiate with valuable properties, it is not the safety providing the value—the value is explained by

²¹ For a related discussion, see Kornblith (2002).

²² This idea was suggested by J. Adam Carter in conversation.

²³ See, for example, Riggs (2009).

the other features, such as the virtuous agency, good environment, rationality etc. In general a property is not valuable merely because it is co-extensive with, corresponds to, supervenes on, arises from, is grounded in, or is caused by, other valuable properties. And so even if safety obtains when and only when these other properties obtain, this does not establish that safety as such is valuable.

A related strategy available to the safety theorist is to claim that safety is a marker that flags the presence of the other properties. According to this strategy, it is not co-extensiveness *per se* that grounds safety's value, but the fact that safety flags the presence of valuable other properties, such as a friendly epistemic environment or good belief-forming process. The value of safety is a kind of evidential value. If the value of safety is value-as-a-marker, however, then it is not valuable in itself, but only derivatively valuable insofar as it helps us attain what is valuable. Such instrumental value is swamped by the presence of the ultimately valuable thing: If the value of safety resides in its instrumental ability to flag friendly epistemic environments, virtuous belief forming methods, and ultimately truth, then whenever we have those things the value of safety is swamped. Thus this suggestion cannot explain a value that safety confers on a true belief.

Safety is not a directly detectable property, furthermore. It is inferred from the presence of other properties, such as friendly environment, virtuous belief forming methods, truth etc., so safety cannot be a 'marker' for these things. We cannot use judgements about a belief's safety to flag whether the epistemic environment is friendly, for example, because we judge that the belief is safe *via* judgements about the friendliness of the environment.

Finally, being a 'marker' requires an internalist condition to obtain in addition to the safety of the belief, namely the internalist condition of the agent's recognising the belief as safe and thereby inferring that other valuable properties are present. This internalist condition is antithetical to the externalist spirit of modal epistemology, and the additional internalist condition required for this proposal locates the purported value of safety outside of the safety condition itself.

Third Strategy: Probabilities of Further True Beliefs

A third strategy to establish the value of safety is to note that if a belief is true and safe it is more likely that the agent will have more true beliefs in the future, solely as a matter of probability, than if the belief is merely true but unsafe. Given that having more true beliefs in the future is valuable, a safe belief is more valuable than an unsafe one.²⁴

This strategy is distinct from the first because it is not committed to the belief's safety *enabling* the agent to form more true beliefs. Instead it posits only that a safe belief is better than an unsafe belief in virtue of the fact that more true beliefs will likely result, as a matter of mere statistical probability, when a person has safe true beliefs compared to a person with merely true beliefs.

We should recall, however, that the truth of a safety counterfactual is not independent from other facts. It is grounded in, and supervenes on, the nature of the epistemic environment, agent, available evidence etc. It is changes in these kinds of actual properties that alter the safety of a belief, and normally safety arises from these properties. With this in mind, it becomes clear that it is not the safety of a belief as such which makes future true beliefs likely, but the good environment, virtuous agency, available evidence etc., which both

²⁴ See also Meylan (2007).

make the belief safe and make future true beliefs likely. To claim a belief is valuable because of its safety is to confuse the order of explanation. Epistemic properties such as good environment and virtuous agency are what cause more true beliefs to obtain in the future, give rise to the safety of the belief, and are plausibly the loci of value. Safety as such does not make likely further true beliefs.

If the person's belief is safe, in other words, it is because other properties are present, and it is these properties that cause the increased likelihood of future true beliefs, and ground the belief's safety. So the correlation is not owed to the safety itself, but only to what the safety supervenes on. It is not clear why standing in a 'common cause' relationship with valuable features renders safety valuable.

Fourth Strategy: What We Care About

Matthew Weiner argues that whether a property of a belief is valuable depends on what we want (our standpoint).²⁵ We use beliefs as premises in practical reasoning, argues Weiner, and whether they are acceptable premises depends on whether we care about, for example, things turning out right for us, being immune from criticism, tracking truth in counterfactual circumstances etc. Weiner argues that if the property of the belief makes it an acceptable premise relative to that standpoint, then it is a valuable property. This proposal suggests a way for modal epistemology to vindicate the primary and secondary value claims: If it is true that we (at least sometimes) care about getting it right in counterfactual circumstances then the safety of a belief will be valuable from that standpoint.

But although it is plausible if we cared about tracking truth in counterfactual circumstances then this could help establish that safety has value, it would be peculiar to care about counterfactual truth tracking as such. We care about being correct in the actual world, and are often ignorant about which the actual world is from a range of epistemically possible worlds, but plausibly we only care about whether we are right in this one. There is arguably no standpoint from which we care about the counterfactual success of our beliefs as such, since counterfactual properties concern only what happens in non-actual worlds. We may care about factors such as getting to truth the right way, not being Gettierised, believing in accord with the evidence etc., and when these things obtain we typically have a safe belief, but having safe belief as such is not itself a goal of enquiry.

We can see this by examining what makes our beliefs safe. Safety and reliability can both be understood subjunctively, but the truth conditions of claims about reliability depend on this world over different times, whereas (according to the Lewisian understanding of counterfactuals) truth conditions of claims about safety concern different possible worlds. We may well care about ourselves in this world over time—this is a sensible thing to care about—but plausibly there is not reason to care about what happens to other possible selves. This is because either they are fictions, and so not a suitable candidate for our concern, or they are concrete but causally isolated from us and so not something that our epistemic situation has bearing on.²⁶ And so Weiner's proposal about the value of knowledge, even if plausible, does not aid the safety theorist in vindicating the value of safety.

²⁵ Weiner (2009).

²⁶ According to modal realism, for example, all counterparts exist, the only thing that can change is which worlds are nearer to which. But which possible worlds are closer per se seems like an odd thing to care about. See also Kripke (1972), p. 45, fn.13.

Conclusion

It does not seem plausible that safety confers value on a true belief. Each attempt described above to explain an unswamped value of safety has failed. Thus it seems that if a belief is true, its being safe as such confers no unswamped value. This suggests that robust anti-luck epistemology, which holds that knowledge is safe true belief, cannot vindicate even the weaker primary value claim: that knowledge is more valuable than mere true belief. And weak anti-luck epistemology can vindicate the weaker claim, since another condition on knowledge might confer value on true belief. But it cannot satisfy the secondary value claim: that knowledge is more valuable than a proper subset of its parts. This is because, if weak anti-luck epistemology is true, any whole comprising all the parts of knowledge other than safety will have the same value as knowledge.

The burden of proof falls to the modal epistemologist who maintains that knowledge is more valuable than any proper subset of its parts to explain the value conferred by a true belief's being safe. Supposing they cannot, we should deem this claim false: "What gives knowledge its distinctive value over true belief relates to its immunity to luck."²⁷

This has consequences for the project of understanding knowledge. We must either eschew the modal conditions as fundamental constituents of knowledge, or maintain that knowledge is not more valuable than that which falls short of knowledge.²⁸

References

- Baumann, Peter (2008). 'Is Knowledge Safe?' *American Philosophical Quarterly* 45(1): 19–30.
- Comesaña, Juan (2005). 'Unsafe Knowledge' *Synthese* 146(3): 395–404.
- Craig, Edward (1990). *Knowledge and the State of Nature*, Oxford University Press.
- DeRose, Keith (2010). 'Insensitivity is Back, Baby!' *Philosophical Perspectives* 24(1):161–187.
- (1995). 'Solving the Skeptical Problem' *The Philosophical Review* 104: 1–52.
- Gettier, Edmund (1963). 'Is Justified True Belief Knowledge?' *Analysis* (23): 121–23.
- Goldman, Alvin and Olsson, Eric (2009). 'Reliabilism and the Value of Knowledge' in A. Haddock, A. Millar, & D. H. Pritchard (eds.) (2009).
- Goddard, Cliff (2010). 'Universals and Variation in the Lexicon of Mental State Concepts' in Barbara Malt and Phillip Wolff (eds.) *Words and the Mind: How Words Capture Human Experience* pp. 72–92.
- Greco, John (2009). "The Value of Knowledge" in A. Haddock, A. Millar, & D. H. Pritchard (eds.) (2009).
- Haddock, A., A. Millar, & D. H. Pritchard (eds.) (2009). *Epistemic Value*, Oxford University Press.

²⁷ Pritchard (2007), p. 99.

²⁸ Many thanks to Bob Beddor, David Black, Adam Carter, Jon Garthoff, Simon Goldstein, Jesper Kallestrup, Peter Milne, Andrew Moon, Duncan Pritchard, Ernest Sosa, and the anonymous reviewers at the *American Philosophical Quarterly* for comments on this essay. Thanks also to audiences at Rutgers University, the Joint Meeting of the South Carolina Society for Philosophy and the North Carolina Philosophical Society, and the British Undergraduate Philosophy Society Conference for helpful discussion.

- Hazlett, Allan (2010). 'The Myth of Factive Verbs' *Philosophy and Phenomenological Research* 80(3): 497–522.
- Kornblith, Hilary (2002). *Knowledge and its Place in Nature*, Oxford University Press.
- Kripke, Saul (2011). 'Nozick on Knowledge' in *Philosophical Troubles: Collected Papers*, Volume 1, Oxford University Press.
- (1972). *Naming and Necessity*, Harvard University Press.
- Kvanvig, Jonathan (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press.
- Lewis, David (1996). 'Elusive Knowledge' *Australasian Journal of Philosophy* 74(4): 549–567.
- (1973). *Counterfactuals*, Harvard University Press.
- Meylan, Anne (2007). 'The Value of Knowledge: Against a Reliabilist Solution' *Proceedings of the Latin Meeting in Analytic Philosophy*.
- Moore, G. E (1903). *Principia Ethica*, Cambridge University Press.
- Nozick, Robert (1981). *Philosophical Explanations*, Cambridge University Press.
- Pritchard, Duncan (2012). 'Anti-Luck Virtue Epistemology' *Journal of Philosophy* (109): 247–79
- (2009a). 'The Value of Knowledge' *Harvard Review of Philosophy* (16): 2–19.
- (2009b). 'Safety-Based Epistemology: Whither Now?' *Journal of Philosophical Research* (34): 33–45.
- (2007). 'Recent Work Epistemic Value' *American Philosophical Quarterly* 44(2): 85–110.
- (2005). *Epistemic Luck*, Oxford University Press.
- Pritchard, D., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford University Press.
- Radford, Colin (1966). 'Knowledge—By Examples' *Analysis* (27): 1–11.
- Riggs, Wayne (2009). 'Understanding, Knowledge, and the Meno Requirement' in A. Haddock, A. Millar, & D. H. Pritchard (eds.) (2009).
- Roush, Sherrilyn (2005). *Tracking Truth: Knowledge, Evidence, and Science*, Oxford University Press.
- Sosa, Ernest (2000). 'Skepticism and Contextualism' *Philosophical Issues* 10(1): 1–18.
- (1999). 'How to Defeat Opposition to Moore' *Philosophical Perspectives* (13): 141–54.
- Swinburne, Richard (1999). *Providence and the Problem of Evil*, Oxford University Press.
- (2000). *Epistemic Justification*, Oxford University Press.
- Vogel, Jonathan (1987). 'Tracking, Closure, and Inductive Knowledge' in S. Luper-Foy (ed.), *The Possibility of Knowledge*, Rowman and Littlefield.
- Weiner, Matthew (2009). 'Does Knowledge Matter?' in A. Haddock, A. Millar, & D. H. Pritchard (eds.) (2009).
- Williamson, Timothy (2000). *Knowledge and Its Limits*, Oxford University Press.
- Zagzebski, Linda (1999). 'What is Knowledge?' in J. Greco & E. Sosa (eds.), *Epistemology*, Blackwell, pp. 92-116.