

Against Swamping

J. Adam Carter • Benjamin Jarvis

(Forthcoming in *Analysis*)

Abstract:

The *Swamping Argument* – highlighted by Kvanvig (2003; 2010) – purports to show that the epistemic value of truth will always *swamp* the epistemic value of any non-factive epistemic properties (e.g. justification) so that these properties can *never* add any epistemic value to an already-true belief. Consequently (and counter-intuitively), knowledge is never more epistemically valuable than mere true belief. We show that the Swamping Argument fails. Parity of reasoning yields the disastrous conclusion that non-factive epistemic properties – mostly saliently *justification* – are *never* epistemically valuable properties of a belief. We close by diagnosing why philosophers have been mistakenly attracted to the argument.

Keywords: knowledge; epistemic value; swamping problem, value of knowledge

1. The Swamping Problem

Recent work by Jonathan Kvanvig (2003), Ward Jones (1997), Richard Swinburne (1999), Linda Zagzebski (2003), Duncan Pritchard (2009a) and others illuminates a particular problem for theorists of knowledge. The problem begins with the widely-held assumption that knowledge is more *epistemically valuable* (henceforth “e-valuable”) than mere true belief.¹ This assumption gives rise to an *ex ante* constraint on the correct theory of knowledge: whatever conditions it places on knowledge (over and above the possession of true belief), a true belief that satisfies all of those conditions must turn out to be more e-valuable than a true belief that doesn’t.² The problem is how to accommodate this *ex ante* constraint.

Failing to solve this problem provides the ground for an objection. Typically, this objection is levied against reliabilist theories of knowledge,³ e.g. one according to which there is only one condition for knowledge over and above the possession of true belief: the condition that (put roughly) the belief must be *reliably formed*. However, the objection has wider scope. Any theorist who is to maintain the *ex ante* constraint must have some plausible explanation for why the e-value

¹ As Kvanvig (2003) notes, the insight that knowledge is to be prized above mere true opinion (even if the former is no practically more useful than the latter) has its roots to Plato’s *Meno*.

² Cf. Baehr (2009) for an argument against construing the value insight as an *ex ante* constraint on an analysis of knowledge.

³ See here Goldman (1992).

of knowledge is greater than that of merely true belief. However the *Swamping Argument* allegedly shows that the value of knowledge collapses into the value of merely true belief.

Our (novel) version of the Swamping Argument rests on three independently plausible premises. A theorist could try to avoid the conclusion by abandoning one of the premises, so we can usefully think of the *Swamping Problem* as a *quadrilemma*: one of the four theses that we will now discuss must be abandoned.⁴

Think of the first, S1, as the value claim that motivates the *ex ante* constraint itself: that knowledge is more e-valuable than mere true belief.

(S1): Knowledge is more e-valuable than mere true belief.

Think of the second, S2, as a specification of the conditions under which a given belief state or process is e-valuable. A familiar position here individuates *epistemic* value from other values, e.g. aesthetic value, moral value. This individuation appeals to a fundamental epistemic good, *truth*; something counts as e-valuable, by way of its connection to that good:

(S2): Any e-value conferred on a belief merely by that belief having some non-factive property is instrumental value relative to the further epistemic good of true belief.

It is by reference to S2 that one might explain why justification is e-valuable; because beliefs that are justified are, in the *default* case, true or alternatively *more likely* to be true than ones that aren't. S2 squares with the idea that we want our beliefs to be justified because we want them to be true.

The third thesis gives a necessary condition for (S1) to be true. Say that a property is a *component* of knowledge if and only if necessarily, a belief that is knowledge has that property. The

⁴ Pritchard (2009a) construes the Swamping Problem as a trilemma. However, our version of the Swamping Thesis is exactly *his* version of the Swamping Thesis.

idea is that knowledge couldn't be more e-valuable than mere true belief unless it has a component other than truth that contributes e-value:

(S3): If knowledge is more e-valuable than mere true belief, then there is a non-factive component of knowledge that, in instances of knowledge, adds e-value.

This thesis might be motivated by the thought that knowledge entails justification, and the additional e-value that knowledge has above and beyond the value of merely true belief is inextricably tied to (although perhaps not wholly or even partly explained by) the e-value that this justification brings with it in instances of knowledge.⁵

The fourth and final thesis is, following here Pritchard 2009a, what we may call the *Swamping Thesis*. This thesis claims:

(S4) If the value of a property possessed by an item is only instrumental value relative to a further good and that good is already present in that item, then it can confer no additional value.

We'll suggest some motivation for the Swamping Thesis in §2 before attacking the thesis in §§3-4.

With the four theses in place, we can construct the Swamping Argument. Let V_e be a function that inputs properties of a belief, and outputs the e-value contributed by those properties. Let K be the property of being knowledge, and T the property of being true. Let C be some arbitrary non-factive component of knowledge. We can reason as follows:

(i) C adds only instrumental e-value relative to the further epistemic good of truth. (by S2)

⁵ It is a virtue of this thesis that it has *none* of the following as consequences: (1) knowledge is factorable into its components (see Williamson (2000)), (2) the e-value of knowledge is explained by the e-value of its components (rather than *vice-versa*), and (3) the e-value of knowledge is the additive sum of the e-values of some of its components (taken individually).

(ii) Therefore, $V_e(T \& C) = V_e(T)$. (by S4)

(iii) Therefore, no non-factive component adds e-value in instances of K . (by Universal Generalization)

(iv) Therefore, $V_e(K) = V_e(T)$ (by S3)

The conclusion (iv) is simply the rejection of S1. If a theorist wants to maintain S1, she must reject S2, S3, or S4. But, each are taken to be independently plausible and they jointly entail that S1 is false, it is natural to conclude that knowledge is not more e-valuable than mere true belief. This is Kvanvig's (2003) ultimate assessment (even though it is specifically a reliabilist theory of knowledge that he has most often cited as an example of a theory that succumbs to the Swamping Problem).⁶ Alternatively, one might reject Kvanvig's revisionist position and take versions of the Swamping Argument as a sort of challenge: a challenge to demonstrate how, on one's own theory of knowledge, the conclusion can be avoided.⁷

Surprisingly, the Swamping Thesis itself appears to have gone almost entirely unchallenged. We think that this is a mistake. In §2, we will advance an argument that shows how the Swamping Thesis has, in conjunction with an equally plausible thesis and S2, the apparently absurd implication that *neither* justification *nor* any other non-factive property of beliefs *ever* confers additional e-value to beliefs. In §§3-4, we diagnose this problem; we explain why the Swamping Thesis should be rejected.

2. More Swamping

In this section, we will show that, S2 and the Swamping Thesis, combined with an equally plausible Swamping Thesis Complement, jointly entail that no non-factive property of belief *ever*

⁶ For Kvanvig, it is another epistemic state, understanding, that bears a sort of distinctive e-value. Of course, this thesis too might be problematic. See §2.

⁷ Cf. Goldman & Olsson (2009) and Olsson (2007) for reliabilist defences against the Swamping Argument. For virtue-theoretic defences, see Greco (2009) and Sosa (2007).

adds any e-value. It is an obvious consequence that no non-factive justificatory component of knowledge adds any e-value—not only in instances of knowledge, but *ever*.

We contend that anyone who accepts the Swamping Thesis (S4) should also accept, *The Swamping Thesis Complement*:

S4*: If the value of a property possessed by an item is only instrumental value relative to a further good, and that good has already failed to be present in that item, then it can confer no additional value.

The kind of analogical reasoning that typically motivates S4 also motivates S4*. For instance, consider a case of a lottery where the winning ticket (drawn at noon on Friday) wins a boat. To make the example clean, suppose only one ticket wins. Let b-value be value related to winning the boat. On Thursday, our sole ticket has some positive instrumental b-value because having it is a means to winning the boat. On Friday, if we show the winning ticket, then, in accordance with S4, our sole (winning) ticket now adds no positive b-value. The idea underlying the Swamping Thesis: our already won boat in conjunction with the ticket that won us that boat is no more b-valuable than the already won boat itself.

However, this knife cuts both ways. Suppose on Friday, that we show the losing ticket. Now, in accordance with S4*, it appears that we should say that our sole (losing) ticket adds no positive b-value. The idea underlying the Swamping Thesis Complement – and arrived at by parity of reasoning – is that the already lost boat in conjunction with the ticket that lost us that boat is no more b-valuable than the already lost boat itself. That is to say, the conjunction is not positively b-valuable at all.

Because the reasoning for each is parallel, we fail to see why someone would accept S4 and not S4*. Notice, however, that together they lead to S4**:

S4**: If the value of a property possessed by an item is only instrumental value relative to a further good, and that good is already present or has already failed to be present in that item, then it can confer no additional value.

Put crudely and metaphorically, the gist of S4** is supposed to be this: after the lottery transaction is over (win or lose) we can throw our ticket out with the trash.

Of course, this suggestion makes trouble in the epistemic case where the good in question is truth. This is because beliefs are apparently more akin to lottery tickets that have *already won or lost* than to tickets for which we are *awaiting a drawing*. At least typically, beliefs are already true or not true.

At this point, we are in a position to construct our main argument. Let *P* be an arbitrary non-factive property of belief. The main argument is as follows:

- (1) Any e-value conferred on a belief merely by that belief having *P* is instrumental value relative to the further epistemic good of true belief. (by S2)
- (2) Beliefs with the property *P* are either already true or already not true.⁸ (Assumption)
- (3) Therefore, *P* confers no additional e-value to these beliefs. (by S4**)
- (4) Therefore, no non-factive property of belief confers any additional e-value to any of the beliefs that have it. (by Universal Generalization)

In other words, no non-factive property—including any sort of non-factive justification—is e-valuable.

⁸ We recognise that there is some scope for objection to (2) as it is currently formulated by appeal to future contingents, which plausibly are not either already true or already not true. However, even if we concede this point, we could still argue to the (still problematic) conclusion that a non-factive property could only be an e-valuable property of future-contingent beliefs.

This result already sounds ridiculous, but arguably, the situation is much worse. The Swamping Argument from §1 can be generalized to show, for any factive property of beliefs F , that F is not more e-valuable than true belief. To do so, we replace S3 with S3*:

(S3*): If F is more e-valuable than mere true belief, then there is a non-factive component of F that, in instances of F , adds e-value.

One can then replace ‘ K ’ with ‘ F ’ in our version of the Swamping Argument in order to conclude that the e-value of F is identical to that of being true.

Putting these results together, we get the apparently untenable conclusion that only factive properties of a belief are e-valuable, and only because they are factive. In other words, the conclusion is that, other than the epistemic good of truth, nothing confers e-value on a belief. This is not a happy conclusion.

Our response will ultimately be to reject both the Swamping Thesis and the Swamping Thesis Complement. However, in §3 we consider other possible responses.

3.Possible Responses

One possible response is to just accept the unhappy conclusion that, as far as e-value is concerned, there is nothing e-valuable *per se* but truth. We don’t think this response merits serious consideration. Our only comment is that, even if this were the right response, it would only show that e-value is entirely uninteresting for epistemology. It would not show, for instance, that knowledge is not worthy of study.

Another possible response would be to deny any version of S3. This would not prevent the first (apparently silly) result that no non-factive property is e-valuable (which does not rely on S3); we will put this response aside as well.

The only possible response that we will take seriously in this section is denying S2, which we repeat below for the reader's convenience:

(S2): Any e-value conferred on a belief merely by that belief having some non-factive property is instrumental value relative to the further epistemic good of true belief.

We will now argue that this response does not offer any real solution.

Suppose some other property besides being true – let's say being knowledge – were taken to be the fundamental epistemic good. Notice that one could substitute an analogous principle for S2 as follows:

(S2 replacement): Any e-value conferred on a belief merely by that belief having some property that is not knowledge-entailing is instrumental value relative to the further epistemic good of knowledge.

Now, letting P now be a property of belief that isn't knowledge-entailing, one could subsequently replace our main argument in §2 with this one:

(5) Any e-value conferred on a belief merely by that belief having P is instrumental e-value relative to the further epistemic good of knowledge. (by S2 replacement)

(6) Beliefs with the property P are either already knowledge or already not knowledge.

(Assumption)

(7) Therefore, P confers no additional e-value to these beliefs. (by S4**)

(8) Therefore, no property of belief that isn't knowledge-entailing confers any additional e-value to the beliefs that have it. (by Universal Generalization)

This result is just as absurd as the prior result that no non-factive property is e-valuable. There is a general problem here, the source of which appears to be that S4** simply prevents any property from conferring any genuine *instrumental* e-value to a belief, no matter what the epistemic good is taken to be. Instrumental e-value is always swamped by either the presence or failed presence of the epistemic good.

This problem remains even if we become pluralists about the epistemic good. For instance, suppose that each of true belief and knowledge were taken to be epistemic goods. Thus, e-value is either truth-related, i.e. t-e-value, or knowledge-related, i.e. k-e-value. The main argument (1)-(4) from §2 can be modified to show that all instrumental t-e-value is swamped by either the presence or failed presence of truth. The argument (5)-(8) above can be modified to show that all instrumental k-e-value is swamped by either the presence or failed presence of knowledge. Consequently, once again, there is no genuine *instrumental* e-value of any sort.

In summary, S4** appears to be problematic whether or not one thinks that truth is the sole epistemic good. Consequently, the right response, we think, is to reject S4** – that is to say, the Swamping Thesis and the Swamping Thesis Complement – whether or not one is inclined to reject S2.

4. A Misguided Analogy

The suggestion that it is the Swamping Thesis and its complement that we ought to reject might sound puzzling in light of the apparent analogical motivation for these theses, e.g. the lottery example from §2. For that reason, we return to the lottery example in order to show that the alleged analogy is problematic.

We think that, in fact, a true belief is *not* analogous to a winning lottery ticket that has been shown so that the prize – in our example, the boat – is now *secured*. Having a true belief is, if

anything, more like holding a winning lottery ticket *after the drawing*, but *before showing it to gain the prize* in such a way that the boat is now in one's possession. Before the winning lottery ticket has been shown and the prize is secured, the winning lottery ticket continues to have instrumental b-value. Indeed, it continues to have instrumental b-value even if one knows that the lottery ticket is a winning ticket. One has to complete a transaction in order to take final possession of the boat before the winning lottery ticket loses its instrumental b-value.

This line of thought leads us to think that a justified true belief may well continue to have instrumental e-value *indefinitely*. Although it is a 'winning' belief, there is never a point at which the belief is 'shown' and the prize is now 'secured' because the belief 'transaction' is now completed. It is useful to remember that beliefs are ongoing *states*, not events with a past terminus.⁹ Just as the work of maintaining a clean house, the work of properly managing a belief is never over. One continues to have to manage the belief properly long after its acquisition; one might very well reflect on it today, and ultimately give it up tomorrow. 'Winning' for a belief is something ongoing rather than something that is, at some point, completed. Consequently, there is nothing obviously absurd about thinking that there continues to be instrumental e-value even when the epistemic good is already present.

Nor is there anything obviously absurd about thinking that there continues to be instrumental e-value even when the epistemic good has already failed to be present. A person with a justified false belief might be taking appropriate means to the epistemic good of truth even if they continue to fall short of possessing it, and moreover, could not possess it by those means. This is not very different from us holding a lottery ticket for the boat when, unbeknownst to us, the lottery is over and our ticket has not been drawn. It would be silly of us to throw our ticket in the garbage under the condition that, *as a matter of fact*, it is a losing ticket if we haven't recognized the condition as such (because, say, we weren't at the drawing and are waiting to hear to the news). It only makes sense to

⁹ Chrisman (forthcoming) and Williamson (2000), 35.

throw out the ticket *because* it is a losing ticket, which implies that we do recognize the ticket *as losing*. So long as the loss is not yet clear to us, the ticket continues to have instrumental b-value for us. In the same way, so long as a person has not gained some further reason to suspect that their belief is not true, there is nothing obviously wrong with thinking that the belief has instrumental e-value for them even if it is, in fact, not true. For example, suppose Serena previously believed her neighbor was home because she saw the front light on. However, she then learned (conclusively) that the neighbor has moved to the Hamptons, but leaves the light on permanently to deter intruders. Thus, Serena is not justified anymore (given what she knows) in believing the neighbor is home. We contend that it is entirely plausible to think that Serena's previous belief was instrumentally e-valuable, but would not have continued to be had she (stubbornly and unreasonably) not given it up after her recent discovery.

More generally, our suggestion is that if a false belief is justified, the subject won't be aware that it is false, so it can easily be instrumentally e-valuable. On the other hand, suppose that the subject is aware that the belief is false or might easily be false. Then, the instrumental e-value of justification is not swamped because there is no justification present, the instrumental e-value of which could be swamped. So, the instrumental e-value of justification is *never* (as opposed to *always*) swamped by falsity.

5. Concluding Remarks

In this paper, our principal suggestion is that the Swamping Problem is, in fact, a pseudo problem: the Swamping Argument fails to provide any reason for worrying that knowledge might not be more e-valuable than mere true belief. (To be clear: we concede that there might be other problems concerning the e-value of knowledge; we simply contend that they are not related to swamping.)¹⁰

¹⁰ Consider Pritchard's (2009b) discussion of the secondary and tertiary value problems.

The Swamping Problem arises primarily from the Swamping Thesis. However, closer inspection casts serious doubt on the Swamping Thesis. We see no more reason to believe in swamping by the presence of the good than to believe in swamping by the failure of that presence. However, believing in both kinds of swamping leads to the conclusion that there is no instrumental e-value no matter what one takes the epistemic good to be). This conclusion is not plausible. We have also suggested that analogical motivation for the Swamping Thesis is poor. It is worth remembering that, even if provocative, analogical reasoning is notoriously fallacious.

References

- [1] Baehr, J. 2009. Is there a value problem? In *Epistemic Value*, ed. A. Haddock, A. Millar, and D. Pritchard, 42-59. Oxford: Oxford University Press.
- [2] Chrisman, M. Forthcoming. The normative evaluation of belief and the aspectual classification of belief and knowledge attributions. *Journal of Philosophy*.
- [3] Goldman, A. 1992. *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, Mass.: MIT Press.
- [4] Goldman, A. And E. Olsson. Reliabilism and the value of knowledge. In *The Value of Knowledge*, ed. A. Haddock, A. Millar, and D. Pritchard, pp. 19-42. Oxford: Oxford University Press.
- [5] Greco, J. 2009. The Value Problem. In *Epistemic Value*, ed. A. Haddock, A. Millar, and D. Pritchard, 313-322. Oxford: Oxford University Press.
- [6] Jones, W. 1997. Why do we value knowledge? *American Philosophical Quarterly* 34(4): 423-439.
- [7] Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- [8] Kvanvig, J. 2010. The swamping problem redux: pith and gist. In *Social Epistemology*, ed. A. Haddock, A. Millar, and D. Pritchard, 89-111. Oxford: Oxford University Press.
- [9] Olsson, E. 2007. Reliabilism, stability, and the value of knowledge. *American Philosophical Quarterly* 44(4): 343-355.
- [10] Pritchard, D. 2009a. What is the swamping problem? In *Reasons for Belief*, ed., A. Reisner & A. Steglich-Petersen, 244-260. Cambridge: Cambridge University Press.
- [11] Pritchard, D. 2009b. The value of knowledge. *The Harvard Review of Philosophy* 16: 2-19.
- [12] Rabinowicz, W. and T. Rønnow-Rasmussen. 2000. A distinction in value: intrinsic and for its own sake. *Proceedings of the Aristotelian Society* 100(1): 33-51.
- [13] Sosa, E. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Vol. I*. Oxford: Oxford University Press.
- [14] Swinburne, R. 1999. *Providence and the Problem of Evil*. Oxford: Oxford University Press.
- [15] Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- [16] Zagzebski, L. 2003. The search for the source of epistemic good. *Metaphilosophy* 34: 12-28.