

# In Defence of Swamping\*

Julien Dutant

12th May 2016

## Abstract

The Swamping Problem shows that two claims are incompatible: (a) that knowledge has more epistemic value than mere true belief and (b) a strict variant of the idea that all epistemic value is truth or instrumental on truth. Most current solutions reject (b). Carter and Jarvis (2012) and Carter et al. (2013) object instead to a principle that underlies the problem. This paper argues that their objections fail and the problem stands. It also outlines a novel solution which rejects (a). By carefully distinguishing value from expected value, one can argue that the greater value of knowledge is merely apparent (Dutant, 2012; Petersen, 2013).

Knowing that something is so seems better than merely to have a true belief that it is (Plato *Meno*, 98a). Hence many accept *Meno's Thesis*, the claim that knowledge

---

\*appeared as Julien Dutant (2013) "In Defence of Swamping", *Thought* 2:4, 357–66. DOI 10.1002/tht3.98.

has more epistemic value than mere true belief.<sup>1</sup> Many are also attracted by the view that anything of epistemic value is either truth — having true beliefs, not having false beliefs — or instrumental on truth. The view is variously called “Epistemic Instrumentalism” (Jones, 1997, 423), “Epistemic value T-Monism” (Pritchard, 2011, 245), “Veritism” (Goldman and Olsson, 2009, 24); call it *Truth Consequentialism about Epistemic Value*.

The Swamping problem purports to show that the two claims are incompatible. It has been used to argue that knowledge is not the highest epistemic good, to reject Truth Consequentialism, and to reject accounts of knowledge that appear wedded to it (see notably Jones, 1997, 426; Zagzebski, 2003; Swinburne, 1999, 58; Riggs, 2002; Kvanvig, 2003, 44–9; Goldman and Olsson, 2009; Pritchard, 2011; Greco, 2010). It relies on a principle against double counting for instrumental value (Goldman and Olsson, 2009, 26–7). In a couple of recent papers, Carter and Jarvis (2012) and Carter et al. (2013) object to the principle, claiming that it faces counterexamples and has absurd consequences. They conclude that the Swamping problem is illusory.

This paper does two things. First, it rebuts Carter, Jarvis and Rubin’s attack on the principle. The counterexamples are flawed, the consequences only follow under controversial assumptions and they are not absurd. So the Swamping problem stands. Second, it highlights an overlooked solution (Dutant, 2012; Petersen,

---

<sup>1</sup>What “*epistemic value*” exactly means — a kind of value, alongside moral value? value that things have conditional on some aim, *e.g.* believing the truth? the value that “epistemic” things such as justification or knowledge tend to have? — is a vexed issue which can be left aside here. See Sosa (2007, 68–9), Pritchard (2011, 244–5), Petersen (2013, 1174–6) for some discussion.

2013). The solution explains away Meno's Thesis as resting on a confusion between value and apparent value. So contrary to what Carter and Jarvis (2012, 695) think, rejecting Meno's Thesis is an option that merits serious consideration.

## 1 The Swamping Thesis

Carter, Jarvis and Rubin target the *Swamping Thesis* from Pritchard's (2011, 248) version of the Swamping problem:

(ST) 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.

The statement has some distracting features: value is ascribed to both properties and items and also identified with properties (a "good" that is "present in an item"), "already" introduces a implicit time index and "additional" an implicit comparison. One way to make it simpler is to value state of affairs — an item's  $a$  being  $F$  is good — and to make the comparison explicit:

(ST\*) If the value of  $a$ 's being  $F$  is only instrumental relative to  $a$ 's being  $G$  then  $a$ 's being  $F \& G$  has no more value than  $a$ 's being  $G$ .

Time is left out but will be brought back in due course. Pritchard (2011, 248), Carter and Jarvis (2012, 692) and Carter et al. (2013, 252) obviously treat (ST) as a schema meant to hold for suitable substitutes of "value" such as "epistemic value".

The thesis is motivated by examples such as the following (Carter and Jarvis, 2012, 693). You own a ticket in a lottery to be drawn on Friday. The prize is a boat and your ticket is the lucky one. On Thursday, owning the ticket has some instrumental value relative to the good of owning a boat. But on Friday, owning the ticket adds no boat-value to owning the boat itself: that is, owning the boat *and the ticket* has no more value relative to the good of owning a boat than simply owning the boat. You can throw the ticket away. Thus if owning the ticket has only instrumental value relative to the further good of owning the boat and you already own the boat, then owning the ticket yields no additional value. Adding the instrumental value of the ticket to that of the boat would be like counting the value of the boat twice (Goldman and Olsson, 2009, 26–7). (ST) prohibits that kind of double counting.

Carter and Jarvis (2012, 693–4) point out that the motivation generalizes. Suppose your ticket turns out to be a loser. On Friday, owning the ticket adds no boat-value to failing to own the boat. That is, not owning the boat *but* owning the ticket is not better than not owning the boat. Hence:

(ST\*\*) If the value of  $a$ 's being  $F$  is only instrumental relative to  $a$ 's being  $G$   
then  $a$ 's being  $F \& G$  and  $a$ 's being  $F \& \neg G$  have no more value than  $a$ 's  
being  $G$  and  $a$ 's being  $\neg G$ , respectively.

A property whose value is only instrumental on  $G$  cannot add value to something in which  $G$  is “already present” or has “already failed to be present”, as Carter and Jarvis put it.

Carter, Jarvis and Rubin object to (ST) with three counterexamples and a *reductio*. First, even after lottery has been drawn, your winning ticket remains instrumentally valuable until it is shown and the prize secured (Carter and Jarvis, 2012, 697). Second, a losing ticket continues to have instrumental value for you even after the draw, as long as you do not know that it is a loser (Carter and Jarvis, 2012, 697). Third, maintaining one’s home well has only instrumental value relative to the good of having a pleasant home. Yet all things being equal, a pleasant and well-maintained home is better than a pleasant one. So being well-maintained confers instrumental value relative to the good of having a pleasant home despite the good being “already present” in it (Carter et al., 2013, 253).

The *reductio* is that (ST\*\*) collapses instrumental value into fundamental value (Carter and Jarvis, 2012, 694–6). Stripped to its essentials, the argument is this. Let  $G$  be the fundamental good for some type of value, say, epistemic value. First, assume  $G$  *Consequentialism about Epistemic Value*:

(GCEV) All epistemic value is  $G$  or instrumental on  $G$ .

Let  $F$  be any property of an item. (GCEV) appears to entail that the item’s being  $F$  has epistemic value only instrumental on  $G$ . By (ST\*\*), the item’s being  $F \& G$  (or  $F \& \neg G$ , as the case may be) has no more value than its being  $G$  (resp.,  $\neg G$ ). Generalizing over  $F$ , no state of affairs adds value to the item’s having or lacking fundamental value. Hence nothing has “genuine” instrumental value. (Carter and Jarvis, 2012, 696) further suggest a stronger *reductio* that does not assume Consequentialism.<sup>2</sup> Whether or not Consequentialism holds for a given type of

---

<sup>2</sup>Thanks to Benjamin Jarvis for pointing this out to me.

value, it clearly holds for *instrumental* value of any type. *Instrumental* epistemic value, for instance, is either fundamental epistemic value or instrumental on it. By the same argument as before, no *instrumental* epistemic value is distinct from fundamental epistemic value; hence instrumental value collapses into fundamental value.<sup>3</sup>

Carter et al. (2013, 253) diagnose the problem as follows: the Swamping Thesis holds only for fundamental goods that “have a clear terminus”. Acquiring a boat has a clear terminus: it is done when you acquire the boat. Swamping occurs “only when the process [of acquiring a boat] is over”. By contrast, having a pleasant home “does not have a clear terminus”. For the home to continue being pleasant, one has to keep on maintaining it well. Hence maintaining it well can “confer instrumental value on the home indefinitely”. So the Swamping Thesis holds only for fundamental goods that “have a clear terminus”. It is doubtful that the diagnosis is correct. While *acquiring the boat* may have “a clear terminus”, *owning the boat* does not. So (ST) should fail when applied value instrumental to owning the boat. But it does not: owning the boat *and the ticket* is no better than owning the boat.

---

<sup>3</sup>Carter and Jarvis (2012, 694n8) are willing to restrict their *reductio* to cases where *a*'s being *G* does not depend on the future. When it does, they allow that *a* is neither “already” *G* nor “already” not-*G*. This merely highlights the need to explicit the time references implied by the locution “being already present in”. Doing so will reveal that the *reductio* requires strong and controversial assumptions.

## 2 The counterexamples

The first purported counterexample confuses the relevant good. The ticket is supposed to have instrumental value relative to the good of *owning a boat*, not relative to the (alleged) good of *owning a winning ticket*. The latter is “already present” as soon as the lottery is drawn, but it is not a good; the former is a good, but it is not “already present” before you acquire the boat. Either way, the first case is not a counterexample to the Swamping thesis.

The second fails to distinguish *value* proper from *apparent value*. Suppose a disgruntled employee is under the impression that sticking pins in a doll is an efficient way to hurt their boss. Does sticking pins in the doll have instrumental value relative to the good of hurting their boss? If the pins do not work, no. It merely *appears* to the employee to have instrumental value. Similarly, if the employee is under the impression that only every other pin works but that they cannot tell which, then each pin appears to them to *have some chance* of having instrumental value. A hallmark of appearances is their dependence on perspectives; that is exactly what we find here. A colleague who knows that the pins are useless would correctly judge that sticking pins lacks instrumental value relative to that good. When the employee themselves discovers their inefficiency, it would be right for them to think that the pins never had any instrumental value. Merely apparent value vanishes with enlightened perspectives.

As long as you do not know the lottery results, your ticket appears to you to have some chance of winning. In that sense “it has some instrumental value for

you” (Carter and Jarvis, 2012, 697). But that is the misleading sense in which we say that the pins are efficient for those who believe that they are: namely, it *appears* to you to have instrumental value. A tell-tale sign that we are dealing with an appearance is the role of perspective. The alleged value is present only as long as you are ignorant of the results. And if somebody informed you, you would not complain that they have thereby destroyed the value of your ticket; you would rather thank them for learning that it had none.

The third counterexample equivocates by leaving times implicit. Well-maintaining one’s home *today* has instrumental value for it being pleasant *tomorrow*. Even though the good of the home being pleasant *today* is “already present” in it, its being pleasant *tomorrow* is not. So we do not have a case where the home is made better by something having merely instrumental value relative to some good “already present” in it.

The point is a general one. A plausible Swamping thesis is relativized to *time-specific instances of goods*. A sterile chicken has less chicken-value than a fecund one, despite the good of being a chicken being “already present” in both. Why? Because the fecund chicken will provide *further instances* of the good of being a chicken (see also Goldman and Olsson, 2009, 26–7). Similarly, a short-lived chicken has less chicken-value than a long-lived one, despite the good of being a chicken being “already present” in both. For the longer-lived chicken provides *further time-specific instances* of the good of being a chicken.

The Swamping Thesis should be specified accordingly. It should also be generalized to cases where an item’s being in a certain way is instrumental to *another*



item having value. Call this the *Explicit Swamping Thesis*:

(EST) If  $a$ 's being  $F$  at  $t$  only has instrumental value relative to  $b$ 's being  $G$  at  $t'$ ,  
then the state of affairs of  $a$ 's being  $F$  at  $t$  and  $b$ 's being  $G$  (resp.,  $\neg G$ ) at  $t'$   
has no more value than the state of affairs of  $b$ 's being  $G$  (resp.,  $\neg G$ ) at  $t'$ .

Times can be instants or longer periods. In the latter case, " $a$ 's being  $F$  at  $t$ " should be read as " $a$ 's being  $F$  throughout  $t$ ". The schema is restricted to cases where  $a$  is  $F$  throughout  $t$  and  $b$  is  $G$  or fails to be  $G$  throughout  $t'$ . The value of an item at a time is the value of its having all the properties it has at that time.

To illustrate, let  $a$  be one's home,  $F$  well maintaining,  $G$  pleasantness and  $t$  today. The counterexample requires a time  $t^*$  such that (i)  $a$ 's being  $F$  at  $t$  only has instrumental value relative to  $a$ 's being  $G$  at  $t^*$  but (ii)  $a$ 's being  $F$  at  $t$  and  $G$  at  $t^*$  is better than  $a$ 's simply being  $G$  at  $t^*$ . If  $t^*$  is today, (i) fails: good maintenance is not instrumental to pleasantness today. If  $t^*$  is tomorrow, (ii) fails: one's home being pleasant tomorrow by having been well-maintained today is no better than its being pleasant tomorrow through other means.

We can also explain why  $a$  itself has greater value today by being  $F \& G$  rather than  $\neg F \& G$ . While its being  $F$  today has merely instrumental value on its being  $G$  tomorrow and is swamped by the latter, the value of  $a$  at  $t$  depends only on its states today and its being  $G$  today does not swap its being  $F$  today.

Suppose by contrast that one's home is going to be destroyed tomorrow. Then its being pleasant and well-maintained today is no better than its merely being pleasant. Yet Carter, Jarvis and Rubin are committed to the view that it is.<sup>4</sup> Of

---

<sup>4</sup>As they stress in Carter et al. (2013, 255n12), their view, as opposed to Olsson's (2007), does

course, it may be reasonable to maintain it well if we are unaware of the impending destruction. But that is merely apparent value.

### 3 The *reductio*

Let us restate the attempted *reductio* with the explicit Swamping thesis (compare Carter and Jarvis, 2012, 696). Let  $G$  be the fundamental epistemic good and  $a$  a belief with property  $F$  distinct from  $G$ :

- (1) The epistemic value of  $F$  is instrumental on  $G$ . (By GCEV)
- (2) Either  $a$  is  $G$  at  $t$  or  $a$  is not  $G$  at  $t$ .
- (3) If  $a$ 's being  $F$  at  $t$  only has instrumental epistemic value relative to  $a$ 's being  $G$  at  $t$ , then  $a$ 's being  $F \& G$  (resp.,  $F \& \neg G$ ) at  $t$  has no more epistemic value than  $a$ 's being  $G$  (resp.,  $\neg G$ ) at  $t$ . (By EST)
- (4) If  $a$  is  $G$  (resp.,  $\neg G$ ) at  $t$  and  $a$ 's being  $F \& G$  (resp.,  $F \& \neg G$ ) at  $t$  has no more epistemic value than  $a$ 's being  $G$  (resp.,  $\neg G$ ) at  $t$ , then  $F$  confers no additional epistemic value to  $a$  at  $t$ . (Definition of "conferring additional value")
- (5) No property distinct from  $G$  ever confers additional epistemic value to a belief. (By 1, 2, 3, 4, generalizing over  $a, F, t$ .)

---

not make the added instrumental value contingent upon the belief having a future. The analogue claim is that good maintenance confers value to one's home even if it is about to be destroyed.

The argument is not valid: there is a gap between (1) and the antecedent of (3).

To close it we need:

- (6) If the epistemic value of  $F$  is instrumental on  $G$ , then  $a$ 's being  $F$  at  $t$  only has instrumental epistemic value relative to  $a$ 's being  $G$  at  $t$ .

The conclusion then follows. Assuming it is absurd, must one blame the Swamping Thesis? No; for one may reject Consequentialism instead. While Carter and Jarvis (2012, 695–6) mention alternative views on the fundamental epistemic good, they do not mention non-consequentialist views of *derived* epistemic goods. Yet many endorse such views. Some think that being *recognizably* good is itself good (Jones, 1997); others that *choosing what is likely to be good* is itself good (Piller, 2009b); others that *indicating the good* is itself good (Olsson, 2009, 102); yet others that earning *credit* for the good is itself good (Riggs, 2002; Zagzebski, 2003; Pritchard, 2010; Greco, 2010). All deny that epistemic value must be either fundamental or instrumental and hence reject (1).

As we mentioned, however, Carter and Jarvis may sidestep the debate by focusing on *instrumental* epistemic value. They replace (1) with the uncontroversial:

- (1\*) The *instrumental* epistemic value of  $F$  is instrumental on  $G$ .

And (6) with:

- (6\*) If the instrumental epistemic value of  $F$  is instrumental on  $G$ , then  $a$ 's being  $F$  at  $t$  only has instrumental epistemic value relative to  $a$ 's being  $G$  at  $t$ .

With the other premises adjusted accordingly, it follows that no property other than the fundamental epistemic good confers instrumental epistemic value. This Carter and Jarvis (2012, 696) take to be absurd.

Assuming it is, must one blame the Swamping Thesis? No; for one may reject (6\*) instead. Non-epistemic analogues clearly fail. Assuming that money only has instrumental value relative to happiness, it is false that somebody's having money now only has instrumental value for *that person's* being happy *right now*. In the epistemic case, (6\*) is rejected by *Future-Belief Consequentialists* who allow a belief to have instrumental epistemic value as a means to its being good *in the future* (Plato, *Meno* 98a; Olsson, 2007; Roush, 2010) and by *Further-Beliefs Consequentialists* who allow a belief to have instrumental epistemic value as a means of *further* good beliefs (Armstrong, 1973, 173; Swinburne, 1999, 64; Goldman and Olsson, 2009, 27–31; see also Olsson's 2011, 876 reading of Kant). For them the *reductio* tells against (6\*), not the Swamping thesis.<sup>5</sup>

These liberal forms of Consequentialism about Epistemic Value face objections, however. Further-Beliefs Consequentialism faces Firth's problem (Firth, 1998, 259–60; Berker, 2013). If a demon rewards you with four hundred good beliefs whenever you form a false belief about pelicans, your false beliefs about pelicans are an efficient means to acquire further good beliefs. That makes them

---

<sup>5</sup>Could this be what Carter, Jarvis and Rubin intend? After all, (ST) may be read as encompassing (6\*) and they do insist on justification being instrumental for the future sustenance of good belief (Carter et al., 2013, 255–6). So perhaps they meant to reject (6\*) only. But if that was so, they could not claim that “the Swamping Thesis itself appears to have gone almost entirely unchallenged” (Carter and Jarvis, 2012, 693) — since the six philosophers mentioned above would count as rejecting it — nor that their view avoids making the instrumental value of justification contingent on a belief's having a future (Carter et al., 2013, 255n12).

pragmatically good relative to an epistemic aim but — arguably — not epistemically good. In Berker’s terminology, Further-Belief Consequentialism violates the “separateness of propositions”: properties of a belief have instrumental epistemic value only insofar as they promote the epistemic value of *that very belief*. Future-Belief Consequentialism faces the Stubbornness problem — at least when Truth is taken to be the fundamental good. A stubborn true belief persists more than a non-stubborn one, yet that does not make the first one better. One may diagnose the problem by imposing a “separateness of times”: properties of a belief at a time have instrumental epistemic value only insofar as they promote the truth of that belief *at that time*.

If one is moved by those objections one will be pushed towards (6\*). With (EST) it follows that no property of a belief confers it *instrumental* epistemic value. Is that absurd? No. Firth’s problem appears to show that being a means to epistemically good things does not make something epistemically good. One may well take that to show that there is no instrumental epistemic value. Note that the view is compatible with there being instrumental value in non-epistemic domains and with there being non-instrumental derived value in the epistemic domain. The limited consequence is far from absurd and not sufficient to reject the Swamping thesis.

## 4 The apparent value solution

The Swamping problem stands. Consequentialism about epistemic value (1) restricted to consequences that the properties of a belief have for *that* belief at *that* time (6) are incompatible with Meno's Thesis. The usual solutions are to adopt more liberal forms of Consequentialism or to reject it altogether. Another solution is to reject Meno's Thesis, for the option "merits serious consideration" (*pace* Carter and Jarvis, 2012, 695).

To see this, consider first a case due to Piller (2009a, 122-3). A lottery has been drawn but the result is not yet announced. You can freely obtain one of two piles of tickets, a large one and a small one. As it happens, the small pile contains the winning ticket. Which pile is better? The small one: it is the means to secure the prize. Which pile should you choose? The big one — in at least one salient sense of "should". Why? Because *in view of the information you have*, it *appears* to have a higher chance of being a means to secure the prize. That we are dealing with an appearance here is suggested by sensitivity to perspectives: to an enlightened onlooker, that choice will not appear likely to be a means to the prize. Note, however, that no mistake in evaluation need be involved. To choose the big pile you do not need to convince yourself that the big pile is good. You need only think that, as far as you can tell, it is likely to be good.

Since we do not perfectly know what is good, we have to go by what appears likely to be good. That can be acknowledged without claiming either that going

by what appears likely to be good is itself good,<sup>6</sup> or that it is not really the case that we ought to go by what appears likely to be good. Going by what appears likely to be good can be genuinely normative without being a value.<sup>7</sup>

There are several notions of information to consider, however. Say that something is *subjectively expectably good* if it appears likely to be good in view of what one believes; *reasonably expectably good* if it appears likely to be good in view of what one reasonably believes; and *knowledgeably expectably good* if it appears likely to be good in view of what one knows. The first two yield different flavours of normativity, sometimes called “procedural” and “substantial” rationality. If you unreasonably believe that the small pile contains the winning ticket, there is a sense in which you ought to choose it (it is “procedurally” rational) but also a sense in which you ought not to (it is “substantially” irrational). One may argue that the third delineates a normatively significant difference as well. If you reasonably believe but fail to know that the small pile contains the winning ticket, there is still a sense in which you ought not to take it.<sup>8</sup>

Now suppose that having a true belief about  $p$  is a final good. Then relative to that final good, believing  $p$  is maximally good (henceforth *optimal*) in view of a body of information that contains  $p$ , but not so in view of a body of information that does not contain  $p$ .<sup>9</sup> If so, relative to the final good of believing the truth about  $p$ , believing  $p$  is subjectively expectably optimal if and only if you believe

---

<sup>6</sup>Though some do: see Piller (2009b).

<sup>7</sup>See Kolodny (2005, 557–8) and Parfit (2011, 160–1) for analogous views.

<sup>8</sup>Though see Littlejohn (2009) for an opposing view.

<sup>9</sup>One may refine the proposal by replacing “contains” by “contains or obviously entails”.

$p$ , reasonably expectably optimal if and only if you reasonably believe  $p$ , and knowledgeably expectably optimal if and only if you know  $p$ . On the assumption that knowledgeable expectations delineate a flavour of normativity, we get a sense in which believing  $p$  when you know  $p$  is recommendable in a way that believing  $p$  when you do not know  $p$  is not.

These facts can be used to explain away Meno's Thesis. Ordinary thinking about "good" and "better" is somewhat inconsistent. It is tempting to say that there is something "good" about choosing the big pile of tickets, though strictly speaking there is no value in doing so. The temptation arises from the fact that in choosing it one would choose as one ought to — since it appears likely to be best from one's perspective. Now from the simple assumptions that believing the truth is good and that knowledgeable expectations yield a distinctive ought it follows that there is a sense in which one believes as one ought to if and only if one knows. That would make it tempting to say that believing when one knows is "better" than believing when one does not, even if it really has no greater value.<sup>10</sup>

Though I cannot elaborate the proposal further here, the sketch should be enough to show that rejecting Meno's Thesis merits serious consideration. Two final points should be stressed. First, I have used the proposal to show that even on their limited assumptions, austere Truth Consequentialists can derive something

---

<sup>10</sup>See Dutant (2012) for further elaboration. Petersen (2013, 1178–9) also rejects Meno's Thesis. His proposal is an error-theoretic version of the "autonomization" solution in Goldman and Olsson (2009): things that generally have good consequences come to be (mistakenly) regarded as intrinsically valuable. As stressed earlier, the present view involves no error. There is something genuinely recommendable or genuinely normative about doing what is likely to be good in view of the information you have. The "recommendability" in question is not illusory but it need not involve value proper.



close to Meno's Thesis. But its core idea is compatible with more liberal views of instrumental epistemic value, more generous views on derived value and alternative conceptions of the fundamental epistemic good.<sup>11</sup> Second, the proposal offers a perspective on the normative role of knowledge that differs sharply from the one commonly adopted in the debate over the value of knowledge. Instead of presenting knowledge as one of the goods we aim at, it presents it as a condition for aiming at goods in a certain way — namely, knowledgeably. Some would try to further explain the distinctive sense in which we ought to aim at goods knowledgeably by saying that knowledgeable aiming is itself good (Piller, 2009b). But the further explanation may be neither successful nor needed.<sup>12</sup>

## References

- Armstrong, D. M. (1973). *Belief, Truth and Knowledge*. Cambridge University Press, London.
- Berker, S. (2013). Epistemic teleology and the separateness of propositions. *Philosophical Review*, 122(3):337–393.
- Carter, J. A. and Jarvis, B. (2012). Against swamping. *Analysis*, 72(4):690–699.
- Carter, J. A., Jarvis, B., and Rubin, K. (2013). Knowledge: Value on the cheap. *Australasian Journal of Philosophy*, 91(2):249–263.

---

<sup>11</sup>Thanks to Carrie Jenkins for pressing that point.

<sup>12</sup>I am grateful to the participants in the NIP's second Early Career conference for useful discussion and to Adam Carter, Benjamin Jarvis and Erik Olsson for detailed written comments on an earlier draft of the paper.

- Dutant, J. (2012). The value and expected value of knowledge. *Dialogue*, 51(1):141–162.
- Firth, R. (1998). Epistemic merit, intrinsic and instrumental. In Troyer, J., editor, *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*, pages 259–71. Rowman & Littlefield, Lanham, Md.
- Goldman, A. and Olsson, E. (2009). Reliabilism and the value of knowledge. In Pritchard, D., Millar, A., and Haddock, A., editors, *Epistemic Value*, pages 19–41. Oxford University Press.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press, Cambridge, New York.
- Jones, W. (1997). Why do we value knowledge? *American Philosophical Quarterly*, 34(4):423–439.
- Kolodny, N. (2005). Why be rational? *Mind*, 114(455):509–563.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press, Cambridge.
- Littlejohn, C. (2009). Must we act only on what we know? *Journal of Philosophy*, 106(8).
- Olsson, E. J. (2007). Reliabilism, stability, and the value of knowledge. *American Philosophical Quarterly*, 44:343–55.

- Olsson, E. J. (2009). In defence of the conditional probability solution to the swamping problem. *Grazer Philosophische Studien*, 79:93–114.
- Olsson, E. J. (2011). The value of knowledge. *Philosophy Compass*, 6(12):874–883.
- Parfit, D. (2011). *On What Matters*, volume 1. Oxford University Press.
- Petersen, S. (2013). Utilitarian epistemology. *Synthese*, 190:1173–1184.
- Piller, C. (2009a). Reliabilist responses to the value of knowledge problem. *Grazer Philosophische Studien*, 79:121–35.
- Piller, C. (2009b). Valuing knowledge: a deontological approach. *Ethical Theory and Moral Practice*, 12:413–28.
- Pritchard, D. (2010). The value problem for knowledge. In Pritchard, D., Millar, A., and Haddock, A., editors, *The Nature and Value of Knowledge*, pages 5–24. Oxford University Press.
- Pritchard, D. (2011). What is the swamping problem? In Reisner, A. and Steglich-Petersen, A., editors, *Reasons for Belief*, pages 244–60. Cambridge University Press, Cambridge, New York.
- Riggs, W. D. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64(1):79–96.
- Roush, S. (2010). The value of knowledge and the pursuit of survival. *Metaphilosophy*, 41(3):255–278.

Sosa, E. (2007). *A Virtue Epistemology*. Oxford University Press.

Swinburne, R. (1999). *Providence and the Problem of Evil*. Oxford University Press.

Zagzebski, L. (2003). The search for the source of epistemic good. *Metaphilosophy*, 34(1-2):12–28.