

Knowledge and Its Value: Some Suggestions as To Why Knowledge Is More Valuable Than True Belief

Simon D'Alfonso

April 9, 2010

Abstract

This paper concerns questions of value in regards to truth and knowledge. Beginning with a cursory outline of the value of truth, I then go on to investigate the value of knowledge in terms of a few specific questions collectively known as the *value problem*. These questions are prompted by consideration of the commonly held view that knowledge is of distinctive value. Firstly, knowledge seems to be more valuable than mere true belief. Secondly, given the value importance bestowed upon it, knowledge should have a greater value than anything which falls short of knowledge. Thirdly, does knowledge have not only a greater quantity, but also a different quality of value than whatever falls short of knowledge? I begin my response to the value problem by proffering a technical argument in support of the view that knowledge does have greater instrumental epistemic value than mere true belief. I then go on to consider a general way of explaining why knowledge has not just a greater amount of, but also a different kind of value than whatever falls short of knowledge. I conclude with a brief look at an alternative, revisionary response to the value problem.

Truth is good. To begin with, it has instrumental value, as true belief is a primary factor in successful action. We generally prefer to form and act upon true beliefs over false ones since we are more likely to get what we want. As William James remarked, truth “is the good in the way of belief” [4, p. 42]. On top of this, it is reasonable to think that truth is more than instrumentally valuable. Most of us at one time or another have wanted to know the truth of a matter for its own sake, indeed as Michael P. Lynch writes, “curiosity is not always motivated by practical concerns” [5, p. 502]. Even if there were no practical consequences or differences between a belief being true

or false, we would surely prefer it to be true and we would also prefer to not believe in trivial falsehoods [5, p. 502]. Furthermore, it can be said that we have a second-order desire for having true beliefs, in that we not only desire the truth, but desire to desire the truth [5, p. 504]. Given all of this, it is fair to say that true belief is an epistemic good and truth an appropriate object of value.

Since knowledge encapsulates truth¹, it follows that if truth is an appropriate object of value, then so is knowledge. This much is clear. However, it is widely thought that knowledge is of distinctive value, that the value of knowledge is distinct not only from the value of its core constituent true belief, but also from whatever falls short of knowledge. Explaining just why knowledge is distinctly valuable, something termed the *value problem*, is to what I shall now turn.

The term *value problem* is taken from Duncan Pritchard, who uses it in his paper ‘The Value of Knowledge’ [7]. In this paper, he defines three levels of the value problem. The *primary value problem* is to explain why knowledge is more valuable than mere true belief. The *secondary value problem* is to explain why knowledge is more valuable than any proper sub-set of its parts. The *tertiary value problem* is to explain why knowledge has not just a greater amount but also a different kind of value than whatever falls short of knowledge.

To what extent should an account of knowledge be judged based on how it deals with the value problem? Should the value problem be a predominant influential factor in the shaping of an account, or merely an aside consideration? I place the significance of the three value sub-problems in decreasing order. Firstly, a successful account of knowledge should have a response to the primary value problem. Secondly, it is important that a successful account is able to accommodate the secondary value problem. Thirdly, whilst it would not be a bad thing for an account of knowledge to be able to answer the tertiary problem, I am not sure that a promising account should be rejected on the basis that it was not able to accommodate the tertiary value problem. At any rate, as I discuss later on, it is arguably not necessary for knowledge to be pitted in a victorious ‘value battle’ against mere true belief in order to justify our particular focus on knowledge in epistemological theorising. So in addressing this issue, one can seek a *validatory* answer by offering an explanation of why knowledge *is* distinctively valuable or a *revisionary* answer by offering an explanation of why although knowledge *is not* distinctively valuable, we are inclined to think of it as so [8, p. 1].

¹If S knows that p , then p must be true.

Beginning with the *primary value problem*, it is firstly important to note that when arguing for the case that knowledge is more valuable than true belief, we are not committed to saying that knowledge is *always* of greater value than corresponding true belief, for cases are conceivable where the true belief that p is of no lesser value than the knowledge that p . Rather, we are committed to the more modest task of arguing that knowledge is *typically* of greater value than corresponding true belief. In this paper, an answer to the primary value problem is to be sought in terms of instrumental value. Taken in isolation, the mere true belief that p will be of just as much instrumental value as the knowledge that p . After all, as Socrates notes in Plato's *Meno*, a true belief about the correct way to the city of Larissa is of just as much instrumental value as knowledge of the way to Larissa; both, if acted upon, will get one to Larissa. Despite this, I contend that overall (i.e. not just with regards to the goal of a particular true belief) the knowledge that p will generally encompass a greater quantity of instrumental value (in terms of true belief) than the mere true belief that p . In order to discuss this idea in more detail, consideration of the 'swamping problem' serves as a good starting point. A standard explication of this problem, due to Linda Zagzebski is as follows:

Imagine two great cups of coffee identical in every relevant respect - they look the same, taste the same, smell the same, are of the same quantity, and so on. Clearly, we value great cups of coffee. Moreover, given that we value great cups of coffee, it follows that we also value reliable coffee-making machines - i.e., machines which regularly produce good coffee. Notice, however, that once we've got the great coffee, we don't then care whether it was produced by a reliable coffee-making machine. That is, that the great coffee was produced by a reliable coffee-making machine doesn't contribute any additional value to it. In order to see this, note that if one were told that only one of the great identical cups of coffee before one had been produced by a reliable coffee-making machine, this would have no bearing on which cup one preferred; one would still be indifferent on this score. In short, whatever value is conferred on a cup of coffee through being produced by a reliable coffee-making machine, this value is 'swamped' by the value conferred on that coffee in virtue of it being a great cup of coffee [9, p. 3].

Although this example specifically relates to reliabilist theories of knowledge, a more general point can be extracted and applied to other epistemological accounts. The gist is that if a property [like that of being reliably formed for beliefs] is only instrumentally valuable relative to some

further good (e.g., true belief or great coffee), then in cases where the further good in question is acquired, the presence of the instrumentally valuable property confers no further value. So as another example, if justification is only instrumentally valuable relative to true belief, then in cases where true belief is acquired, the justification involved in its formation confers no further value.

It would be rash to conclude that the swamping problem in effect shows that there can be nothing (including knowledge) more epistemically valuable than mere true belief. To get a better hold of the issue, let us begin with a formulation of the swamping problem given by Pritchard, in terms of an inconsistent triad of claims [9, p. 5]:

(1) The epistemic value conferred on a belief by that belief having an epistemic property is instrumental epistemic value relative to the further epistemic good of true belief.

(2) 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.

(3) Knowledge that p is sometimes more epistemically valuable than mere true belief that p .

Now, since it follows from (1) and (2) that (3) is false, these claims are jointly inconsistent. So what modifications could be made to resolve this inconsistency? Claim (2) seems right, therefore I will take it as the first to remain. Claim (3) is the first to question. Whilst (3) is a claim which I would ultimately like to endorse, even if it were to be rejected, there are still other notions of value besides epistemic value which one could appeal to in order to support the claim that all in all knowledge has greater value than true belief.

For instance, knowledge can be said to have greater *practical* value than mere true belief. This type of response is taken up by Plato in the *Meno*, where it is noted that although a true belief about the correct way to Larissa is of just as much instrumental use as knowledge of the way to Larissa, someone who has knowledge about the way to Larissa might have certain practical advantages over someone who only has a true belief about the way to Larissa. For example, if the path starting off on the journey looks dubious, the person without knowledge might be tempted to turn back, whereas the other person will continue down the path certain with the knowledge that it is the correct path [10]. So despite the swamping problem, any greater value of knowledge over mere true belief need not be understood only in terms of instrumental value relative to the good of true belief, for there could be other benefits of having knowledge over mere true belief.

Given this, one option is to accept (1) and (2) whilst replacing (3) with a more general claim about the greater value of knowledge, “that knowledge is generally of greater all things considered value than mere true belief whilst simultaneously granting that from a purely epistemic point of view there is no additional value to be had” [9, p. 7].

But we shouldn’t be too hasty in settling for such a relaxing modification to (3). To begin with, the claim only requires that we explain why knowledge is *sometimes* more epistemically valuable than mere true belief, a less burdensome task than explaining why knowledge is *always* or *often* more epistemically valuable than mere true belief. Also, consider the following situation:

Suppose that someone comes to you and says that in a moment one of two scenarios will obtain: either one will have a true belief that p or one will have knowledge that p (where one does not know which proposition is at issue). Furthermore, it is stipulated that all the practical consequences are kept fixed in both scenarios, so there will be no practical benefit to choosing the one option over the other. Nevertheless, shouldn’t one choose knowledge rather than mere true belief? [9, p. 8]

Unlike those who defend (3), those who deny it outright or settle for a relaxed version as just discussed, seem unable to account for this intuition. So let us make it clear that this investigation into the distinctive value of knowledge is in terms of epistemic value. Predictably this leads to consideration of claim (1). This claim is susceptible to scrutiny by those who can argue that a value conferred on a belief by that belief having an epistemic property can be of a suitable value type other than *instrumental* epistemic value. But prior to considering this type of strategy, I would like to expound a technical argument for the case that knowledge that p is sometimes more epistemically instrumentally valuable than mere true belief that p . I would say that where in this apparently inconsistent triad of claims this argument fits is a leak resulting from the fact that (1) and (2) do not speak sufficiently for knowledge to rule out (3).

Now, a ‘template’ definition of knowledge that identifies the core element of true belief will suffice for the technical argument I have in mind. This simple definition consists of

(1) Truth

(2) Belief

(3) X

where X is a set of 1 or more conditions. The basic idea is that knowledge that p , because of the X component (whether it includes justification, reliability, undefeatability, etc.), is generally of more instrumental epistemic value than mere true belief that p , because it generally encompasses or leads to more true beliefs and less false beliefs, it is ‘truth conducive’.

To help get this point across, I would first like to emphasise that as I see it, the swamping problem is limited in what it covers. Comparing the knowledge that p against the mere true belief that p is not just about comparing the true belief that p which is a constituent of knowledge against the mere true belief that p . Overall, it is about comparing the *knowledge* that p , including all it involves, against the mere true belief that p . Thus, one can argue that there is something, other than the constituent true belief that p , which is involved with the knowledge that p and which makes it, *ceteris parabus*, more instrumentally valuable. Here are some varied examples to help convey this notion. The first example, although simple and by no means definitive, is a good starting point.

Often knowledge of something involves other true beliefs, whereas its corresponding mere true belief simply does not. Here are some examples to help convey this notion. The first is a modest one, but serves as a good entry point. Consider a situation where Bob and John form beliefs about who attended a party. Bob guesses correctly that Tony went to the party. John on the other hand, has the knowledge that Tony went to the party via this deductive process:

- Tony or Terry went to the party.
- Terry did not go to the party.
- Therefore Tony went to the party.

So the knowledge that ‘Tony went to the party’ involves the knowledge that ‘Terry did not go to the party’; in this case, there are two true beliefs associated with the knowledge that ‘Tony went to the party’. Bob on the other hand has one true belief and a 50% chance of guessing that ‘Terry did not go to the party’. Of course, the knowledgeable true belief that p is not necessarily tied to any other related true beliefs as it is in this case. For example, John’s knowledge could have been obtained via a direct process, like him seeing Tony at the party or someone telling him that

Tony went to the party, and in these cases the true belief would stand alone. But the general point remains that at least sometimes, probably more often than not, the knowledge that p involves one or more extra true beliefs on top of the true belief that p . Furthermore, cases along a continuum can be conceived that strengthen this point, leading to examples where the knowledge that p in some sense necessarily requires the knowledge, hence true belief, that q ($Kp \supset Bq$). Take this simple mathematical example. Suppose someone thinks about the question ‘how many factors does the number 152 have?’. If they were to form the mere belief that ‘152 has 6 factors’, their belief would be true, requiring no other true beliefs. If, on the other hand, they wanted to *know* that ‘152 has 6 factors’, they would proceed by calculating the factors of 152. In doing so, it follows that they would also come to have the true belief that ‘1, 2, 4, 38, 76 and 152 are factors of 152’.

The next example hones in on the idea that knowledge, unlike mere true belief, can be used as an instrument in the subsequent generation of other true beliefs. Thus, whilst the mere true belief that p and the knowledge that p might have the same immediate instrumental value, often the instrumental value of knowledge is potentially greater overall.

Bob and John are independently hiking across a forest. They both begin their hike from the same starting location and they both have the same destination. Bob leaves before John and they never meet. The forest has one and only one path that joins their starting location and destination. Along this path there are forks, at which the hikers must choose whether to turn left or right. If they choose correctly, they remain on the right path. If they choose incorrectly, they hit a dead end or find themselves on another path. As it happens, the branching of this path follows a pattern that corresponds to certain features of the environment surrounding the forks. These certain features are caves, streams and boulders.

Now, for each fork, if a cave is located in between the left and right turning paths of the fork then the correct path to take is the right one, unless there is a boulder at the fork but no stream, in which case the correct path to take is the left one. If there is no cave but there is a stream, then the correct path to take is the left one and if there is no cave and no stream then the correct path to take is the right one. This is formally captured with the following, where $\text{Left}(x) \equiv \neg \text{Right}(x)$:

$$\begin{aligned}
&(\forall x)((\text{Fork}(x) \wedge \text{Cave}(x) \wedge \text{Stream}(x) \wedge \text{Boulder}(x)) \supset \text{Right}(x)) \\
&(\forall x)((\text{Fork}(x) \wedge \text{Cave}(x) \wedge \text{Stream}(x) \wedge \neg \text{Boulder}(x)) \supset \text{Right}(x)) \\
&(\forall x)((\text{Fork}(x) \wedge \text{Cave}(x) \wedge \neg \text{Stream}(x) \wedge \text{Boulder}(x)) \supset \text{Left}(x)) \\
&(\forall x)((\text{Fork}(x) \wedge \text{Cave}(x) \wedge \neg \text{Stream}(x) \wedge \neg \text{Boulder}(x)) \supset \text{Right}(x)) \\
&(\forall x)((\text{Fork}(x) \wedge \neg \text{Cave}(x) \wedge \text{Stream}(x) \wedge \text{Boulder}(x)) \supset \text{Left}(x)) \\
&(\forall x)((\text{Fork}(x) \wedge \neg \text{Cave}(x) \wedge \text{Stream}(x) \wedge \neg \text{Boulder}(x)) \supset \text{Left}(x)) \\
&(\forall x)((\text{Fork}(x) \wedge \neg \text{Cave}(x) \wedge \neg \text{Stream}(x) \wedge \text{Boulder}(x)) \supset \text{Right}(x)) \\
&(\forall x)((\text{Fork}(x) \wedge \neg \text{Cave}(x) \wedge \neg \text{Stream}(x) \wedge \neg \text{Boulder}(x)) \supset \text{Right}(x))
\end{aligned}$$

Also, although the streams and boulders can always be seen, the cave entrances can sometimes become swamped by forest growth so that the only way to see a cave and make sure one is present is by rummaging through the growth to find its entrance.

This pattern is not something that Bob and John generally have to worry about, because they rely on trusty maps to guide them. John though, being the cartophile he is, thoroughly studied his map prior to the hike and acquired knowledge of this pattern. Bob on the other hand, gave his map only a cursory glance prior to setting off. Now, although Bob and John both brought maps with them, as misfortune would have it, they both lose their maps, coincidentally around the same area. Although they do not know it, at the point they lose their maps, Bob and John are three correct fork turns away from reaching their destination. The correct turns to make at these three remaining forks are right, right and left respectively; the first fork (fork1) has a cave and a stream, the second fork (fork2) has a cave and a stream and the third and final fork (fork3) has a cave, no stream and a boulder.

Without a map to assist him, Bob makes the assumption that at all forks with a cave the correct turn to make is a right turn. He approaches the first remaining fork, sees a cave and correctly turns right. He approaches the second remaining fork, sees a cave and correctly turns right. Although he forms the true beliefs that the first and second forks require right turns, this is obviously not knowledge because it is based on reasoning involving a false assumption. Bob's reasoning can be formally represented thus:

$$\begin{aligned}
&(\forall x)(\text{Fork}(x) \wedge \text{Cave}(x) \supset \text{Right}(x)) \\
&\text{Fork}(\text{fork1}) \wedge \text{Cave}(\text{fork1}) \\
&\text{Fork}(\text{fork2}) \wedge \text{Cave}(\text{fork2})
\end{aligned}$$

$\therefore \text{Right}(\text{fork1}) \wedge \text{Right}(\text{fork2})$

When he arrives at the third fork however, Bob does not see a cave (because it is swamped by forest growth). Rather than attempting to determine whether or not a cave is present by rummaging through the growth, he uses the one piece of knowledge that he somehow gathered from his cursory glance of the map: the fork following two right-turn forks always has a cave. So using his true belief that the correct first turn was right and his true belief that the second correct turn was right, Bob forms the true belief that there is a cave at the third fork somewhere under the forest growth:

$(\text{Right}(\text{fork1}) \wedge \text{Right}(\text{fork2})) \supset \text{Cave}(\text{fork3})$

$\text{Right}(\text{fork1}) \wedge \text{Right}(\text{fork2})$

$\therefore \text{Cave}(\text{fork3})$

With the true belief that there is a cave at the third fork and his false assumption that all forks with caves require right turns, Bob turns right again, this time incorrectly, and ends up on the wrong path.

The knowledge-based actions of John on the other hand result in a successful outcome. As mentioned, John has complete knowledge of the three features and their correspondence to correct turns. Based on this knowledge, he correctly turns right at both the first and second forks. At the third fork, he observes a boulder but no stream and he does not see a cave. But in order to make the correct decision, he needs to determine whether or not there is a cave amongst the growth. Like Bob he also has and uses the knowledge that the fork following two right-turn forks always has a cave. So this knowledge, coupled with the knowledge that the first two correct turns are right turns, gives John the knowledge that there is a hidden cave somewhere around the third fork. Armed with this knowledge, the knowledge that there is no stream, and the knowledge that there is a boulder, John correctly turns left and reaches his destination.

So Bob uses his mere true beliefs that the first and second forks require right turns and ends up with a false belief. On the other hand, John uses his knowledge that the first and second forks require right turns and ends up with another true belief. Thus when the time comes to act upon their beliefs, the knowledgeable John ultimately succeeds whereas the mere truth believer Bob fails. In the long run, he who has knowledge over mere true belief has the greater true belief generation potential.

With the point of these examples, the epistemic instrumental value (EIV) of the knowledge that p is being assessed in terms of not only the constituent true belief that p , but also any other true beliefs that it encompasses or leads to. Put simply, if the knowledge that p also encompasses or leads to the true belief that q and $\text{EIV}(\text{Belief}(p)) + \text{EIV}(\text{Belief}(q)) > \text{EIV}(\text{Belief}(p))$, the knowledge that p would overall have greater instrumental epistemic value than the mere true belief that p , in virtue of the fact that it generates a greater overall quantity of true belief. Under these terms, whilst the knowledge that p would not be of greater value than the mere true belief that p *plus* the mere true belief that q , such a response to my argument would miss its crux. For if there is possibly going to be a greater total instrumental value (in the way of extra true beliefs) associated with the knowledge of a *single* proposition as opposed to the mere true belief of *that single* proposition, then that is grounds for the claim that knowledge that p is of greater epistemic instrumental value than mere true belief that p , and a reason for preferring the former over the latter. In the example just given, John’s knowledge gives him one extra true belief over mere truth believer Bob.

Furthermore, knowledge is not only going to result in more true beliefs, but also less false beliefs. As a simple example, take John, who is waiting at home for his friend Peter when he hears someone knocking on his front door. He believes that the person at the door is his friend Peter and hence forms the belief that the person at his door has a name that starts with ‘P’. As it turns out, the person at the door does have a name that starts with ‘P’, so his derived belief is true. However, the person who is at the front door is actually another one of John’s friends, namely Paul. Compared with the mere true belief which John has, that the person at his door has a name that starts with ‘P’, if John had had the knowledge of this fact instead, say, because he knew it was Paul due to a certain pattern to the knock that only Paul does, then John would have had two true beliefs instead of one true belief and one false belief.²

Linda Zagzebski writes “in the sense most commonly discussed by reliabilists, truth conduciveness is a function of the number of true beliefs and the proportion of true to false beliefs generated by a process” [13, p. 465]. So couched in these terms, the point here is that the knowledge that p is more *truth conducive* than the mere true belief that p , and truth conduciveness is translatable to epistemic instrumental value.

On a brief final note, there is perhaps another way to look at the epistemic instrumental value of knowledge that is independent of any appeal to the value of constituent true belief. Zagzebski

²This type of example was adapted from Bertrand Russell [11, p. 131]. It is an early example of arguments found in Russell’s work on epistemology that cover Gettier-notions several decades before Gettier’s paper.

continues

There is another sense of truth conduciveness, however, which is important at the frontiers of knowledge and in areas, like philosophy, that generate very few true beliefs, no matter how they are formed. I suggest that we may legitimately call a trait or procedure truth conducive if it is a necessary condition for advancing knowledge in some area even though it generates very few true beliefs and even if a high percentage of the beliefs formed as the result of this trait or procedure are false. [13, p. 465]

This other sense of truth conduciveness can be tied in with another way in which knowledge has greater instrumental epistemic value than mere true belief. Basically, the X component of knowledge can be seen to have this type of truth conduciveness by comparing mere false belief with ‘false knowledge’, where ‘false knowledge’ consists of

- (1) Falsity
- (2) Belief
- (3) X

Once again, here X is a set of 1 or more conditions, such as justification, reliability, undefeatability, etc. Although a mere false belief that p is valueless, the ‘false knowledge’ that p could still be truth conducive in this sense in virtue of X.

Whilst I have just shown that it is unnecessary to appeal to a type of value other than instrumental epistemic value in addressing the primary value problem, the other value problems, particularly the tertiary, prompt consideration of other types of epistemic values which could be associated with knowledge. In [7] Pritchard points out that Zagzebski’s ‘swamping problem’ argument fails to recognise an important value-theoretic distinction, that something can have non-instrumental, extrinsic value. It can have *final value*, where it is valued for its own, non-instrumental sake. For example, consider

a book printed on the first ever printing press and an exact replica produced by lasers. We would undoubtedly value the former more than the latter even though we can accept, for the sake of argument, that these two objects have the same relevant intrinsic

properties. Moreover, we would value the former for its own sake, quite apart from what instrumental value it had. The difference in value, however, is clearly due to the relational properties of the objects concerned. The reason for this difference in value is that a book printed on the first ever printing press, unlike an exact replica, has final value - it is valuable for its own sake because of how it is produced. [7, p. 6]

Thus if another type of epistemic value as such can be found, it could be used to resolve the tertiary and *a fortiori* the secondary value problems. Exploring this avenue, Pritchard considers the applicability of virtue epistemology in approaching a resolution. According to virtue epistemology, knowledge is true belief resulting from epistemic virtue. Virtue-theoretic accounts of knowledge

hold that knowledge is to be understood as true belief that arises out of one's intellectual [epistemic] virtues. There are various ways of understanding the intellectual virtues, some more permissive than others, but the common thread is that such virtues are at least reliable cognitive traits which are stable parts of one's cognitive character. [7, p. 7]

Roughly speaking, virtue theoretic accounts identify knowledge with cognitive achievement and it is this association with cognitive achievement that confers final value on knowledge. However, whilst such virtue-theoretic accounts of knowledge can provide us with the right value resources to address the value problem, they arguably fall short of being adequate accounts of knowledge. Pritchard shows this by providing examples of both knowledge without virtue epistemological achievement and virtue epistemological achievement without knowledge [7, pp. 12-16].

But perhaps, without considering any particular account of knowledge, we can countenance a more general reason for why we might value knowledge over whatever falls short of knowledge, why knowledge has final value beyond the instrumental and intrinsic value of its constituents. Involving notions similar to those involved in the idea of knowledge-as-achievement and final value, it rests on a crucial difference between cases like the swamping problem example, in which the coffee is a product of a coffee machine and cases in which belief is a product. This difference concerns the nature of the products' origins/causes. In the case of the coffee, its cause is of no value significance, for as far as one is concerned the machine is purely and totally an instrument for the production of the product (i.e. the coffee). With beliefs however, the causes, the originators of beliefs are

humans, and although humans are instrumental in belief formation, they are also agents of value judgements and are bound to assess and judge their products (e.g. beliefs) not only intrinsically, but also with regard to extrinsic factors like the viability and quality of their production methods. A small modification to the coffee example will elucidate this point. Say, we replace the coffee making machines with two people. One person, Bob, has never made a coffee in his life. Another person, John, is one of the world's finest baristas. John, as usual, produces an optimal cup of coffee. Bob, as luck has it, throws a bunch of coffee ingredients together and produces a fluke coffee that is just as good as Johns. Despite the instrumental and intrinsic value equivalences of these coffees, I get the feeling that one would opt to consume a coffee put together by the world's best barista rather than a coffee produced by someone who is not associated with the principles of coffee making excellence. The barista-made coffee would have a special final value because of who made it. The only unique value that Bob's 'luck coffee' would have is *shock value*. To further hone in on this point, consider if Bob were to make the effort to go to barista school. Upon graduation, he goes to the effort to gather the finest coffee ingredients and puts together a coffee which is of the same instrumental and intrinsic value as the coffee he randomly produced in the previous example. Surely Bob would nonetheless value this coffee over the 'luck coffee', because of his efforts in its causal history. Likewise, the principles of knowledge are the principles of true belief formation excellence. Hence if someone forms a true belief by following these principles, they are going to value their product more than a corresponding product which does not have these principles and efforts associated with it. These extrinsic features of knowledge, such as the factors in its causal history and its intimate connection with the truth out there, would explain its added extrinsic value.

In closing, I would like to briefly discuss another way to think about the value problem. For those who do not buy such appeals to final value, this might provide a more plausible option. Even if attempts to give a validatory response to the value problem are inadequate or inconclusive, I think that this line of response would at least serve as a safe revisionary response of sorts. I contend that even if knowledge were not distinctively valuable, there is justification for its key place in epistemological theorising. This revisionary response applies to the standard range of positions that treat true belief as a basic epistemic good which has value. It simply appeals to the teleological relationship between value and the deontic, where the deontic is antecedent to and explained in terms of the evaluative.

Since true beliefs have value, it is good to have true beliefs. Since it is good to have true beliefs,

we ought to pursue knowledge, since it is truth conducive and maximises the chances that our beliefs attain their goal of truth and avoid falsity. All of this is summed up in the following claim

If you want to have true beliefs then you ought to have knowledge

This then can explain the intuition that knowledge seems better than mere true belief. Marian David suggests something along similar lines when he writes that these intuitions³ “arise due to a confusion of sorts. They do not reflect any bonus of intrinsic value accruing to knowledge over and above ... [mere] true belief ...; rather they reflect our desire to have our desires satisfied” [2, p. 310]. Thus there is a *normative* difference between true belief (the good) and knowledge (the ought).⁴ In light of this, it is not hard to see why knowledge, an epistemological ought (because it is true conducive), is analysed and sought.

³David lists the following four intuitions, of which only the first is directly relevant to the current discussion

1. Knowledge seems better than mere true belief
2. Justified true belief seems better than unjustified true belief
3. Unjustified false belief seems worse than justified false belief
4. Unjustified true belief versus justified false belief? Intuition hesitates.

⁴This idea is explored in [6]

References

- [1] Baehr, Jason S, ‘Virtue Epistemology’, *The Internet Encyclopedia of Philosophy*, James Fieser and Bradley Dowden (eds.), URL = <<http://www.iep.utm.edu/v/VirtueEp.htm/>>.
- [2] David, Marian, ‘Truth as the Primary Epistemic Goal: A Working Hypothesis’. In Matthias Steup and Ernest Sosa (eds.) *Contemporary Debates in Epistemology*. Wiley-Blackwell, 2003, pp. 296-310.
- [3] Greco, John, ‘Virtue Epistemology’, *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/epistemology-virtue/>>.
- [4] James, William. *Pragmatism and the Meaning of Truth*, Harvard UP, 1975.
- [5] Lynch, Michael P. ‘Minimalism and the Value of Truth’. *The Philosophical Quarterly*, Vol. 54, No. 217, 2004, 497-517.
- [6] Piller, Christian. ‘Valuing Knowledge: A Deontological Approach’. *Ethical Theory and Moral Practice*, Vol. 12, No. 4, 2008, 413-428.
- [7] Pritchard, Duncan. ‘The Value of Knowledge’. *The Harvard Review of Philosophy* (forthcoming) URL = <<http://www.philosophy.stir.ac.uk/postgraduate/documents/ValueOfKnowledge2.pdf>>.
- [8] Pritchard, Duncan. ‘Knowing the Answer, Understanding and Epistemic Value’. *Grazer Philosophische Studien*, Vol. 77, 2008, 325-339.
- [9] Pritchard, Duncan. ‘What is the Swamping Problem?’. In A. Reisner & A. Steglich-Petersen (eds.) *Reasons for Belief*. (Springer, forthcoming). URL = <<http://www.philosophy.ed.ac.uk/staff/documents/WhatIsTheSwampingProblem.pdf>>.
- [10] Pritchard, Duncan, ‘The Value of Knowledge’, *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/knowledge-value/>>.
- [11] Russell, Bertrand. *The Problems of Philosophy*. Oxford University Press, 1997. [Originally Published 1912]
- [12] Schroeder, Mark, ‘Value Theory’, *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/value-theory/>>.

- [13] Zagzebski, Linda. 'Virtues of the Mind'. In Ernest Sosa and Jaegwon Kim (eds.) *Epistemology: An Anthology*. Blackwell, 2000, pp. 457-467.