

Commodious Knowledge

Christoph Kelp, KU Leuven Institute of Philosophy

*Mona Simion,
KU Leuven Institute of Philosophy and Arche Research Centre, University of St Andrews*

Abstract

This paper offers a novel account of the value of knowledge. The account is novel insofar as it advocates a shift in focus from the value of individual items of knowledge to the value of the commodity of knowledge. It is argued that the commodity of knowledge is valuable in at least two ways: (i) in a wide range of areas, knowledge is our way of being in cognitive contact with the world and (ii) for us the good life is a life rich enough in knowledge.

1 Introduction

We care a lot about knowledge. As a society, we invest a lot of time and energy in the development of institutions whose aim it is to accumulate or distribute knowledge. Universities, schools, libraries and the internet are among the most prominent of these. On an individual level, we send our children to school and encourage them to go to university so that they can acquire knowledge about a wide range of topics. Some of us go to considerable financial lengths in order to make this possible. Finally, in philosophy, the study of knowledge has historically received a great deal of attention. A lot of effort has been made to get clear on what exactly is involved in knowing.

From a philosophical point of view, the fact that we seem to care so much about knowledge gives rise to a number of interesting questions: First, is our concern with knowledge warranted? In other words, does knowledge have value that is *special* at least in the sense that it would warrant this concern? Second, in what respect(s) exactly is knowledge valuable?

These two questions are closely related. If knowledge does have the kind of value that would warrant our concern with it, then we may expect there to be a story about the specific respect(s) in which

knowledge is valuable. The first question thus naturally gives rise to the second. On the other hand, an answer to the second question will allow us to make progress towards answering the first question. After all, if we have an account of the respect(s) in which knowledge is valuable, we will be better positioned to assess whether the value that attaches to knowledge warrants our concern with it.

The aim of this paper is to provide novel answers to both of the above mentioned questions. More specifically, we will try to develop a novel way in which knowledge is valuable, thereby answering the second question. In addition, we will argue that, as a result, the first question also receives a positive answer: the novel way in which knowledge is valuable serves to vindicate our concern with it.

2 The Value Problem

2.1 Two Challenges

What does it take to provide a satisfactory answer to the question whether knowledge has value that is special enough to warrant our concern with it? There are a number of proposed answers in the literature. They differ from one another in the strength of the demands imposed.

Let's begin with what is widely regarded as the most lenient proposal, which dates back as far as Plato's *Meno*. To begin with, notice that it is nearly universally accepted that knowledge requires true belief.¹ Now suppose it turns out that knowledge is in no respect more valuable than true belief that isn't knowledge, i.e. mere true belief. In that case it would seem wrong to care about knowledge rather than mere true belief. Our special concern with knowledge would seem misplaced. This motivates a first constraint on satisfactory accounts of the value of knowledge:

[C1] Any satisfactory account of the value of knowledge must explain why knowledge is in some respect more valuable than *mere true belief*.

Some have claimed that simply meeting C1 won't be enough to give a satisfactory account of the value of knowledge. Jonathan Kvanvig, for one, argues that more is needed: Suppose knowledge consists of a set of constituents. Suppose, next, it turns out that knowledge

¹ Notable exceptions are David Lewis [1996] and Colin Radford [1966] who have denied that knowledge requires belief and Allan Hazlett [2010] who has challenged the thesis that knowledge is factive.

is in no respect more valuable than some proper subset of its constituents. In that case it would be wrong to care specifically about knowledge rather than the proper subset of equally valuable constituents. Our special concern with knowledge would still seem misplaced. In view of these considerations, Kvanvig favours the following constraint:

[C2] Any satisfactory account of the value of knowledge must explain why knowledge is in some respect more valuable than *any proper subset of its constituents* [Kvanvig 2003: xii-xiii].

2.2 Riggs's Requirement

It has widely been taken for granted that the task of explaining the value of knowledge consists in showing that individual items of knowledge are more valuable than individual beliefs that fall short of knowledge. Wayne Riggs, for one, is very clear about this when he proposes that the best way to understand C2 is as follows:

$(\forall s)(\forall p)[\text{Value}(sKp) > \text{Value}(sRp)]$ (where R is some relation comprising elements of K , and $R \neq K$)

[Riggs 2009: 334]

Roughly, the idea here is that in order to meet C2 we need to show that it is better for one to know *that* p than to have a belief *that* p that falls short of knowledge, for all propositions p .

While the challenge Riggs unpacks in the above quote is of course C2, it is not hard to see that C1 can be given the same treatment: for all propositions p , knowledge that p must be more valuable than mere true belief that p . Key to this way of fleshing out the challenges is that they require showing that *every item* of knowledge is more valuable than the corresponding belief that falls short of knowledge. (In what follows, we will refer to this requirement as 'Riggs's requirement'.)

Let's get one thing out of the way: Everything we will say in this paper is compatible with the idea that that knowledge is valuable in a way that satisfies Riggs's requirement. By way of explanation, consider first the following distinction between two types of value: final and instrumental.² For something to be *finally* valuable is for it to be

² For more on different types of value see e.g. [Korsgaard 1983; Kagan 1998; Rabinowicz & Rønnow-Rasmussen 2000] all of which are reprinted in [Rønnow-Rasmussen & Zimmermann 2005].

valuable for its own sake. For instance, it is widely acknowledged that happiness is finally valuable. In contrast, for something to have *instrumental* value is for it have value as a means to an end. For instance, it is widely acknowledged that money is instrumentally valuable, at least in the kinds of society we live in: it allows its possessors to buy things that enable them to achieve a certain level of comfort in life.³ The reason everything we say is compatible with the idea that knowledge is valuable in a way that complies with Riggs's requirement is that it is compatible with all our claims that knowledge but not true belief that falls short of knowledge is finally valuable. If so, every item of knowledge is more valuable (and indeed enjoys a different kind of value) than the corresponding true belief that isn't knowledge.

That said, for the purposes of this paper, we'd like to bracket responses to the challenges in terms of the final value of knowledge and focus on responses in terms of instrumental value instead. Curiously, once we are clear that we are aiming for this kind of response, the prospects of success look somewhat dim. The reason for this is that, when focusing solely on instrumental value, some items of knowledge appear to be of no value whatsoever. (Let's call items of knowledge that have no instrumental value whatsoever items of 'useless knowledge'.) Knowing the exact number of grains of sand in the jar you brought back from last year's summer holiday is but one popular example [e.g. Sosa 2003]. If some items of knowledge are useless, however, they are no more valuable than the corresponding beliefs that fall short of knowledge, which are useless as well. A satisfactory response to C1 and C2 appears no longer available.

³ Traditionally, final value was thought to be the same as intrinsic value, i.e. value that accrues to something in virtue of its intrinsic properties only. Moreover, intrinsic value was thought to contrast with instrumental value. However, there has been a growing consensus in recent value theory that final value need not be intrinsic. A famous example of final extrinsic value is that of Princess Diana's dress, which is valuable for its own sake, but derives this value at least in part from its relationship to Princess Diana [Rabinowicz & Rønnow-Rasmussen 2000: 121]. Similarly, some have argued that instrumental value is not the only kind of non-final value. Michael Zimmermann [2015: §6] considers the case of a medical test the results of which indicate that the patient is in good health. Zimmerman notes that it is not implausible to say that, in this case, the results are also good. In contrast with good health, which is arguably finally valuable, the positive results are not valuable for their own sake and hence are not finally valuable. What's more, the good results are not a means to good health either and so aren't instrumentally valuable. As a result, the type of value that attaches to the results has to be non-final non-instrumental value.

On the other hand, recall that C1 and C2 are to be motivated by our concern with knowledge. That is to say, we wanted an account of the value of knowledge that explains our concern with knowledge. If so, however, it is far less evident that a satisfactory account of the value of knowledge must satisfy Riggs's requirement. After all, our concern with knowledge does not appear to extend to all items of knowledge. In particular, we seem to have little concern for items of useless knowledge—and, we want to add, rightly so. If so, examples of useless knowledge are best understood as suggesting that a satisfactory account of the value of knowledge need not satisfy Riggs's requirement.

Once we abandon Riggs's requirement, we might attempt to replace it by a weaker requirement. Rather than venturing to show that each item of knowledge is more valuable than the corresponding mere true belief, we might think that it will do to show that enough items of knowledge are more valuable than the corresponding beliefs that aren't knowledge. We may try to rise to this challenge in at least two ways. First, we may try to explain the value of knowledge in terms of a property (or properties) of *individual items* of knowledge. While we do not mean to deny that this is a promising line to pursue, in what follows we'd like to explore a different approach, which starts from the thesis that knowledge is a commodity and aims to explain the special value of knowledge in terms of a property of *the commodity*.

3 The Value of Knowledge

3.1 Commodity Value

'Knowledge' is a mass term, like 'water'. It is widely agreed, however, that mass terms denote stuff that can be measured but not counted.⁴ In the case of knowledge and water the stuff is a kind

⁴ Pelletier [2010: 128]. Pelletier also points out that mass terms are generally regarded to have divisive and cumulative reference: one can subdivide the stuff to which a mass term can be correctly applied infinitely and the mass terms will continue to be correctly applicable to its parts (divisive reference) and one can add as much of the stuff to an existing quantity of it as one likes and the term will continue to be correctly applicable to the resulting mass (cumulative reference). One might think that this is implausible for knowledge. After all, if we take away the justification component from a bunch of beliefs that make up some quantity of knowledge, for instance, 'knowledge' won't be correctly applicable to the resulting mass. But notice that this is not a special problem for 'knowledge'. After all, if

of commodity—something one can have more or less of.⁵ Now suppose that it can be shown that the commodity of knowledge has special value, that an account of the value of the commodity can be given that satisfies C1 and C2. There is reason to believe that this will also be sufficient to adequately meet these challenges. After all, an account of the value of the commodity would make good sense of our concern with knowledge. What's more, if we did succeed in providing an account of the value of the commodity of knowledge we would still stand as a good chance as any of vindicating the special focus on knowledge in the history of epistemology. For that reason, it seems that everything is to be gained and nothing to be lost by exploring the prospects for an account of the value of the commodity of knowledge.

Before moving on to the value of the commodity of knowledge, we'd like to take a look at the value of another central commodity in our lives, to wit, water. Water is of course valuable in many respects. For the purposes of this paper, we'd like focus on one valuable quality of water, its power to quench our thirst. Suppose liquid hydrogen were just as well suited to quench our thirst as is water. The constituents of liquid hydrogen, H_2 , are a proper subset of the constituents of water, H_2O (so that liquid hydrogen stands to water as, for instance, justified true belief stands to knowledge.) Now suppose you have before you two glasses, one containing water, the other liquid hydrogen. It is plausible that the glass with water is no more valuable than the glass with liquid hydrogen, at least not with respect to its power to quench your thirst. After all, *ex hypothesi*, liquid hydrogen is as well suited to do the job as water is. Does that mean that water, the commodity, couldn't have special value, value that warrants our concern with water? No. To see this, suppose (as happens to be the case) that liquid hydrogen is extremely rare and can exist only in very special environments. Suppose that, at the same time, water is easily available in a wide range of places and to

you chop up some water molecules that make up a quantity of water and take away the oxygen atoms, say, 'water' won't be correctly applicable to the resulting mass either. Pelletier [2010: 129] suggests that this problem may be solved by distinguishing between metaphysical and semantic facts. The thought is that while as a matter of metaphysical fact water consists of hydrogen and oxygen atoms, this is not recognised semantically, at least not in English. Whether or not this solution works need not concern us here. After all, so long as the problem is not exclusive to 'knowledge', we can, for present purposes at least, safely ignore it.

⁵ See [Treanor 2013, 2014] for a view similar to this as well as an illuminating account of how one might measure amounts of knowledge.

a wide range of people. In that case water is plausibly valuable to us in a way that would warrant our concern with it. What makes water thus valuable is not just the fact that it has the power to quench our thirst. After all, we are supposing that water shares this property with liquid hydrogen. It is a combination of the fact that it has this property and the fact that it is so widely and easily available. To put a snappy label to it, water is of special value because it is our way of quenching our thirst.

Now, we want to suggest that the situation is in essence the same with knowledge on the one hand and true belief that falls short of knowledge on the other. One valuable property of knowledge is that it is a way of correctly representing the world around us. It is undeniable that the same holds for true belief that falls short of knowledge. Now compare two agents, *A* and *B*, such that *A* knows that *p* and *B* truly believes but doesn't know that *p*. It is very plausible that *A*'s belief that *p* is no more valuable than *B*'s, at least not with respect to its correctly representing the world—just as it is very plausible, in the imagined case above, that the glass of water is no more valuable with respect to thirst-quenching than the glass of liquid hydrogen. Arguably, however, just as in the case of water and liquid hydrogen, this result is compatible with knowledge being valuable in a way that would warrant our concern with it. In fact, the very same properties that account for the special value of water in the imagined case account for the corresponding special value of knowledge: in a wide range of areas, knowledge is widely and readily available.

To see this, consider first perceptual beliefs about middle-sized dry goods. On any non-sceptical account of knowledge, given formation by suitable processes (alternatively: on suitable grounds) in sufficiently hospitable epistemic environments, these beliefs will qualify as knowledge. For instance, my belief that there is a computer on the desk before me qualifies as knowledge: it is produced by a highly reliable ability to recognise tables in an epistemically hospitable environment. Now the crucial point is that, for beliefs in this range, formation by suitable processes in hospitable environments is the norm; formation of beliefs by unsuitable process, or in inhospitable environments is the exception. If this isn't immediately clear, consider again my belief that there is a computer on the desk before me and ask yourself what would have to be the case for my belief to remain true but fall short of knowledge. Those with some training in epistemology will find it easy to answer this question: I mistake a

hologram for a computer, whilst unbeknownst to me there is a computer somewhere else on the desk, I acquire my belief by a highly unreliable process such as a coin-toss, etc. While any of this might come to pass, it is undeniable that, as a matter of fact, it only rarely does. For that reason, cases of knowledge are the norm and cases of true belief that fall short of knowledge are the exception.

Perceptual beliefs about middle-sized dry goods are not the only cases in point. Consider testimonial belief about propositions of crucial practical importance in our lives: propositions about bills that need to be payed, the nature of the sickness of your child and the medication that will cure it, what's available at the local restaurant, etc. Or consider inferentially supported beliefs that exploit a variety of natural and social regularities: that my car is still parked outside the institute, that Cameron is still the prime minister of the UK, etc. Here too, when beliefs in these ranges are formed by suitable processes in sufficiently hospitable epistemic environments, they will qualify as knowledge. Here too, cases of knowledge are norm and cases of true belief that fall short of knowledge are the exception.

These considerations suggest that, in wide range of cases, knowledge is widely and readily available. All we have to do to acquire knowledge is open our eyes, listen to what other people tell us, attend to our feelings, etc. In comparison, in those areas true belief that falls short of knowledge is a rare commodity that exists only in very special environments. In parallel with the case of water, then, what makes knowledge specially valuable is not just the fact that it involves a correct representation of the world. It is the fact that it has this property in combination with the fact that, in a wide range of areas, it is so widely and easily available. Just as water is of special value because it is our way of quenching our thirst, knowledge is of special value because, in a wide range of areas, it is our way of correctly representing the world.⁶

Before moving on, we'd briefly like to say a few words about what exactly the proposed account of the special value of the commodity of knowledge amounts to. To this end, we'd like to return once more to the parallel account of the special value of the commodity of water. In particular, we'd like to point out that, on the present account,

⁶ One could point out that we tend to attach higher value to scarce commodities than to readily available ones. However, it is not market value that is at stake here; it is instrumental value related to our epistemic goals. Analogously, while the market value of truffles is higher than that of bread, it is undeniable that, due to its ready availability, bread is more instrumentally valuable for our survival.

the surplus value of the commodity of water over the commodity of liquid hydrogen in the society under consideration is *contingent* and *agent-relative*. To see that it is contingent, note that there may be other possible worlds at which liquid hydrogen exists in abundance and water is rare. At those worlds, liquid hydrogen may well be more valuable to members of the counterpart of our society than water. It is agent-relative because there may be creatures inhabiting the world of our society for whom liquid hydrogen is more valuable than water, despite being so rare. One example are creatures who are severely allergic to any liquids containing oxygen. And, of course, the same holds, *mutatis mutandis*, for the value of the commodity of knowledge: on the proposed account, the surplus value of knowledge over true belief that falls short of knowledge is contingent and agent-relative.

3.2 Value Inheritance

Suppose that the above account of the value of the commodity of knowledge is successful. One question that we might ask ourselves at this stage is what, if anything, this entails for the value of individual items of knowledge.

One possible answer is that individual items of knowledge inherit special value from being instances of a specially valuable type and that, in consequence, individual items of knowledge turn out to be more valuable than individual items of true beliefs that fall short of knowledge. (We will henceforth refer to the thesis that tokens of a valuable type inherit value from the type as 'Value Inheritance'.)

Now we think that Value Inheritance is most plausible when the value of the relevant type is *final* value. For instance, if happiness as a type of state is finally valuable, one might think that instantiations of happiness are finally valuable also. At the same time, we said that we would bracket the issue of final value for the purposes of this paper. So the crucial question is whether Value Inheritance holds for instrumental value.

In a recent paper, Alvin Goldman and Erik Olsson [2009] argue for a positive answer. They offer the state of possessing money as a case in point. The idea is that possessing money is an instrumentally valuable type of state as it frequently produces states that are finally valuable (e.g. happiness). Crucially, Goldman and Olsson claim that "*each* token of this type inherits instrumental value from the type" [Goldman & Olsson 2009: 32]. If they are right, Value Inheritance might still hold for instrumentally valuable types in some cases.

It may look as though, compatibly with Value Inheritance, Goldman and Olsson can allow for the existence of token states of possession of money that are useless. As they point out, each token state inherits value from the type even though they may not actually produce finally valuable states. In the case of money, this may happen when the money isn't spent at all or it is badly invested [Goldman & Olsson 2009: 32]. On reflection, however, it cannot be the case that a token state of possessing money is genuinely useless (i.e. has zero instrumental value) and yet inherits positive instrumental value from the instrumentally valuable type. One of the two has to go. The question is which one and why.

Given that Goldman and Olsson accept that Value Inheritance holds for possessing money, it has to be useless money. What appears to be going on here is that the value of token states of possessing money resides in a *dispositional* property—roughly, its power to bring about something finally valuable (though perhaps only by producing something else of instrumental value). Since a dispositional property can be had even when the disposition is not manifested, if the value of a token state of possessing money resides in the dispositional property, money can be valuable even when it does not bring about something of final value. Crucially, however, in that case the token state is not useless.

But couldn't there be useless money? We think there could be. Suppose you own a one hundred Euro bill. Unfortunately, however, the bill is on an unmanned spaceship that is now in such a remote part of the universe that it is certain to have disintegrated before it can reach the next living being. In this case, we submit, the token state of your owning this bill is useless. There is a rationale behind this verdict. Even when the instrumental value of a certain item resides in a dispositional property, the item can be useless if there is no chance at all that the disposition will be manifested. This is what is happening in (a suitably fleshed out version of) the above case. There is simply no longer any chance that the money is spent and hence that the disposition in which the money's value resides will be manifested. That's why the state of your owning the bill in this case is useless. There is thus reason to believe, *pace* Goldman and Olsson, that Value Inheritance does not hold for money.⁷

⁷ Goldman and Olsson also offer a second example, namely that of good motives for actions. The idea is that good motives are a type of state that acquires value from its relation to good action. Over time, good motives have come to be

In view of these considerations, we are suspicious of the idea that Value Inheritance generally holds for instrumental value. This is of course not to say that Value Inheritance doesn't often hold in particular cases. But so long as Value Inheritance does not hold for instrumental value in general, the proposed account of the value of knowledge can allow for the possibility of useless knowledge. We'd also like to emphasise that, while we suspect that Value Inheritance does not generally hold for instrumental value, we are not wedded to this idea. On the contrary, we are open to be convinced otherwise.⁸ Suppose it can be established that Value Inheritance holds for all cases in a certain range and that all cases of knowledge are within that range. All that follows from our account is that individual items of knowledge are more valuable than the corresponding beliefs that aren't knowledge after all. While this would require us to change our verdicts about the value of beliefs about the number of grains of sand in some jar and like cases, the result is not unwelcome. After all, prospects for an alternative explanation of the intuition of zero value are fairly bright: in the case of knowledge, the value of the type is often overlooked and, additionally, instrumental value is not always inherited by tokens of the type. No surprise that the inherited value could remain unrecognised in the case of knowledge.

valued independently, "in themselves" [Goldman & Olsson 2009: 33] as they put it. Now, it is not clear to us what exactly Goldman and Olsson mean here. Is it that good motives come to be valued for their own sake? Whatever the answer to this question may be, we think that it is very plausible that good motives are valued for their own sake. As we already pointed out, we are also sympathetic to idea that Value Inheritance holds when the value of the type is final value. So, we are happy to grant that Goldman and Olsson have succeeded in identifying a case of Value Inheritance. It's just that the case is not of the kind we were looking for. After all, at least for all Goldman and Olsson have argued, the value inherited here is final, not instrumental.

⁸ A promising place to look are normative properties. For instance, possession of legal tender is a type of state that has the instrumentally valuable normative property of entitling you meet financial obligations by using it. We think that the idea that the instrumentally valuable normative property is inherited by each token of the state type carries promise. If, in addition, knowledge has instrumentally valuable normative properties, we would have what it takes to argue that knowledge has instrumental value that is inherited by each token item of knowledge. That said, we will not pursue this line in any more detail here. Suffice it to say that the instrumental value that we argue attaches to the commodity of knowledge is not a normative property. Even if Value Inheritance holds for normative properties, there still is little reason to think that it holds for the properties that, according to our proposal, make the commodity of knowledge valuable.

4 Objections and Replies

With our account of the value of knowledge in play, we'd like to discuss a couple of objections that one might press against our account of the value of knowledge. The first points to a putative disanalogy between the case of water and the case of knowledge. In contrast, the second charges the account of the value of knowledge with being too thin.

4.1 Disanalogy

One worry one might have is that there is a crucial disanalogy between the case of water and the case of knowledge. In particular, one might think that, in the case of water, the following counterfactual is true: "If there hadn't been any water, we wouldn't have been able to quench our thirst, at least not on a global scale." After all, there wouldn't have been enough liquid hydrogen around to do the trick. In contrast, it might be thought that the parallel counterfactual for knowledge is not true: "If there hadn't been any knowledge, we wouldn't have been able to correctly represent the world, at least not on a global scale." After all, there would still have been enough true beliefs around to do the trick. In this respect, then, knowledge and water are disanalogous. What's more, the disanalogy is important. The truth of the counterfactual is just what makes it plausible that, in the thought experiment, water is of special value in the first place. Given that the parallel counterfactual about knowledge isn't true, the proposed account of the value of knowledge fails.

We want to make four points by way of response. First, the closest worlds at which there isn't any knowledge (or at least not enough to satisfy our needs for correct representation on a global scale) are very far off. Knowledge abounds, whilst true belief that isn't knowledge is rare not only at the actual world but a good way out into modal space, well past the boundary even of nearby worlds. That is to say that knowledge is very robustly our way of correctly representing the world. Given that this is so, it is less than clear that, even if the counterfactual is indeed false, this compromises the value of knowledge that, according to our proposal, attaches to knowledge at the actual world.

Second, we are somewhat suspicious of the idea that if there hadn't been any knowledge, there still would have been plenty of true beliefs around, enough to satisfy our need of correctly repre-

senting the world on a global scale. The reason for this is that it is not clear to us that the following counterfactual isn't true instead: if there hadn't been any knowledge, there wouldn't have been any beliefs around. Very roughly, the thought here is (i) that many of our beliefs are formed by processes that were selected for reliably producing true beliefs and (ii) that had these processes not been reliable enough to produce knowledge, we would not have been reliable enough in finding food, avoiding predators and other dangers etc. to make it as a species.

Third, suppose we stipulate that, at the world of our thought experiment, water abounds and liquid hydrogen is rare (S1) and that had water been rare, liquid hydrogen would have abounded (S2) with the result that the allegedly troublesome counterfactual is false for water as well. In that case, water is still of special value at the world of the thought experiment. After all, it is the only resource present at that world suited to quench our thirst on a global scale. At best, the counterfactual shows that something else might have been of special value instead of water (and would have been had water not been of special value). Note, however, that this is entirely unproblematic on the current account as the value of water is acknowledged to be contingent.

Crucially, the case is exactly the same for knowledge and true belief that falls short of knowledge at the actual world. Although knowledge entails true belief and, in consequence, true belief abounds at the actual world, true belief that falls short of knowledge is not entailed by knowledge—in fact, it excludes knowledge. Even though knowledge and true belief are abundant at the actual world, true belief that isn't knowledge is rare. This gives us the analogue of S1. The counterfactual pressed by the objector is the analogue of S2, i.e. that if knowledge had been rare, true belief that isn't knowledge would have abounded. What we can now see is that even if our objector is right and the counterfactual is true, our proposed account of the value of knowledge will still get off the ground. Just as in the case of water the truth of S2 does not impugn the special value of water at the world of the thought experiment, the truth of the parallel counterfactual does nothing to impugn the special value of knowledge at the actual world. After all, just as in the case of water, knowledge is the only resource present at the actual world suited for correctly representing the world on a global scale.

Fourth, the case of water/liquid hydrogen at the world of the

thought experiment and of knowledge/true belief that isn't knowledge at the actual world is not to be confused with the case of oil/alternative energy sources at the actual world. In the latter case, even though we use oil to meet our energy-related needs, we might just as well have used alternative energy sources instead. For that reason, it might be thought, oil is not of special value at the actual world. Even if this assessment of the value of oil at the actual world is correct, there is a crucial difference between this case and the cases of water/liquid hydrogen and knowledge/true belief that isn't knowledge. After all, in contrast with the cases of water and knowledge, in this case, the alternative energy sources also abound at the actual world (alongside oil). This means that oil is not the only resource present at the relevant world suited to meet our energy-related needs. As a result, even if we agree that oil is not of special value at a world at which there are many other resources present suited to do the job of oil, this does not compromise the plausibility of the thesis that water and knowledge are of special value at the relevant worlds in virtue of being the only resources present at the relevant worlds suited to do the relevant jobs.

4.2 Thinness

Another worry one might have is that the proposed account of the special value of knowledge is too thin. Recall that the value of knowledge is both contingent and agent-relative. If that's really all there is to the value of knowledge, do we genuinely capture the special value we have long believed knowledge to have?

Here is our response to this worry. First, it is entirely compatible with the proposed account that there are other ways in which knowledge is valuable. We already mentioned one such alternative, albeit only to put it aside for the purposes of this paper, to wit, that knowledge is finally valuable. However, there may be even further ways in which knowledge may be valuable, including alternative ways of being instrumentally valuable. As a result, it is compatible with our account that the proposed account does not fully capture the value we believe knowledge to have. In that case, it doesn't matter that our account of the value of knowledge is too thin.⁹

Second, suppose it turns out that our thin way is the only way

⁹ Note also that, for water, too, there may be other ways in which it is valuable. For instance, it may be that water but not liquid hydrogen is suitable for the purposes of irrigation.

in which knowledge is valuable over and above true belief that falls short of knowledge. Suppose furthermore that our account does not fully capture the special value we have believed knowledge to have. In that case, we believe that we have to revise our beliefs about how special knowledge is. However, it does not detract from the significance of the account. Even if knowledge turns out to be less special than we thought it to be, it will be of some importance to know that it is in some way better than true belief that falls short of knowledge.¹⁰

Third, and relatedly, it is not so clear to us just how special we believe knowledge to be. It is not even clear to us that there exists a uniform opinion on the issue. For instance, it might turn out that epistemologists have believed knowledge to be much more special than ordinary folk. That said, we do believe that the observation we started from is pretty solid: as a society we care about knowledge in the sense that we invest a lot of time and energy in institutions aiming at the accumulation or distribution of knowledge and so on. Even if our account is too thin to capture the special value some (or even all) of us may have believed knowledge to have, it will be of considerable significance whether or not knowledge is valuable in a way that vindicates this concern. Crucially, there is reason to think that our account delivers on this front. To see this, let's return one more time to the case of water. Suppose the society under consideration cares about water in much the same way in which we care about knowledge: it invests a lot of time and energy in institutions investigating its nature, as well as institutions involved with the management of water, its sourcing, distribution, quality control and so on. Suppose it turns out that the only thing that makes water more valuable than liquid hydrogen is that water is a more widely and readily available commodity than liquid hydrogen. Even so, this result will vindicate the society's concern with water. By way of evidence, note that it makes sense for them to continue caring about water in the way they have done: to invest time and energy in institutions investigating its nature and to be concerned with its management. Moreover, it would not make sense to significantly alter the way they have cared for water. In particular, it would not make sense for them to stop caring more for water than for liquid hydrogen, to allocate fifty percent of their resources to research into the nature and management of liquid hydrogen and so on. In this way, there is reason to believe that the result concerning the value of water vindicates the society's concern

¹⁰ Again, it is easy to see that the same may hold in the case of water.

with water. And the same holds, *mutatis mutandis*, for our concern with knowledge.

5 The Superiority of Knowledge

5.1 A Third Challenge

In §2 we outlined two challenges that any satisfactory account of the value of knowledge will have to meet. While these challenges enjoy rather widespread support, some think that there is a further condition of adequacy on accounts of the value of knowledge. In particular, Duncan Pritchard argues that a satisfactory account of the value of knowledge must, in addition, show that knowledge enjoys a different kind of value than true belief that falls short of knowledge. According to Pritchard, then,

[C3*] Any satisfactory account of the value of knowledge must explain why knowledge is in some respect more valuable than any proper subset of its constituents *not just as a matter of degree but as a matter of kind* [Pritchard *et al.* 2010: 7-8].

On the account developed in §3, knowledge is valuable because it is our way of correctly representing the world. Note that this will serve to address the challenges from §2 (C1, C2) as it explains why knowledge is more valuable than mere true belief and more valuable than beliefs that fall short of knowledge. At the same time, it is not hard to see that our account does precious little to address C3*, according to which knowledge must have a different kind of value than beliefs that fall short of knowledge. For that reason, we'd now like to take a closer look at Pritchard's challenge.

First, we'd like to express a worry about the challenge. Suppose it can be shown that knowledge is finally valuable. In that case, it would seem that we have everything we could hope for. By the same token, any plausible challenge for an account of the value of knowledge should at this stage be met. Unfortunately, there is no guarantee that C3* will be met. To see this, suppose that true belief turns out to be finally valuable as well. Given, additionally, that all other kinds of value are equally shared between true belief and knowledge, there will no kind of value that attaches to knowledge that does not attach to true belief that isn't knowledge. C3* will not be met. As a result, there is reason to believe that C3* is too demanding.

At the same time, we think that Pritchard may have been on to something when he introduced C3*. To see this, let's take a look at how he motivates it:

[I]f one regards knowledge as being more valuable than that which falls short of knowledge merely as a matter of degree rather than kind, then this has the effect of putting knowledge on a kind of continuum of value with regard to the epistemic, albeit further up the continuum than anything that falls short of knowledge. The problem with this 'continuum' account of the value of knowledge, however, is that it fails to explain why the long history of epistemological discussion has focused specifically on the stage in this continuum of value that knowledge marks rather than some other stage (such as a stage just before the one marked out by knowledge, or just after). Accordingly, it seems that accounting for our intuitions about the value of knowledge requires us to offer an explanation of why knowledge has not just a greater *degree* but also a different *kind* of value than whatever falls short of knowledge.

[Pritchard *et al.* 2010: 7-8]

What becomes clear here is that Pritchard takes it, first, that no account of the value of knowledge on which it is on a continuum with the value of true belief that isn't knowledge can be successful. He also seems to think, second, that the only way in which we can avoid placing knowledge on such a continuum is by showing that knowledge enjoys a different kind of value.

Importantly, the second claim is false. Even if a difference in kind of value is sufficient to get knowledge off the value continuum with true belief that falls short of knowledge, it isn't necessary. There are other ways in which one type of good, A, can be discontinuous in value with another type of good, B. For instance, it may be (i) that any amount of A is better than any amount of B (henceforth also 'Strong Superiority') or (ii) that some amount of A is better than any amount of B (henceforth also 'Weak Superiority'). Mill famously put forth such a discontinuous account of value relations:

It is quite compatible with the principle of utility to recognise the fact, that some kinds of pleasure are more desirable and valuable than others. – Of two pleasures, if [...] one of the two is, by those who are competently acquainted with both, placed so far above the other that they [...] would not resign it for any quantity of the other pleasure which their nature is capable

of, we are justified in ascribing to the preferred enjoyment a superiority in quality, so far outweighing quantity as to render it, in comparison, of small account.

[Mill 1963: 210]

Both Strong and Weak Superiority will take knowledge off a value continuum with true belief that isn't knowledge. At the same time, neither requires a difference in kind between these two.

It comes to light that, even though C3* is arguably too demanding, Pritchard's motivations for C3* may still justify a condition of adequacy for accounts of the value of knowledge over and above C1 and C2, to wit:

[C3] Any satisfactory account of the value of knowledge must show that knowledge is not on a value continuum with true belief that falls short of knowledge.

Crucially, even if showing that knowledge enjoys a different kind of value than true belief that falls short of knowledge is sufficient to meet C3, it is not necessary. Rather, to rise to C3 it will be enough to show that knowledge is, in some respect, at least weakly superior to true belief that falls short of knowledge. That is to say, we need to show, at a minimum, that some amount of knowledge is in some respect better than any amount of true belief.

5.2 Eudaimonic Value

While we aren't certain that even C3 constitutes a reasonable demand on adequate accounts of the value of knowledge, in the remainder of this section we want to try and provide some support for the claim that knowledge is weakly superior to true belief that isn't knowledge. We take this part of the paper to be rather speculative: a sketch of an argument that it may be worth pursuing in more detail elsewhere rather than a thorough defence. Roughly, the idea is that a certain amount of knowledge is required to achieve one of the highest goods in life: human flourishing or what Aristotle called 'eudaimonia'. Eudaimonia is a type of happiness. Crucially, however, it is not happiness of any old sort. As Rosalind Hursthouse points out, eudaimonia is "the sort of happiness worth seeking or having." [2007: §2]

We will assume that a eudaimonic life is at least weakly superior to a life without eudaimonia. To put it in Mills's terms, no one fully acquainted with both lives would sacrifice a eudaimonic life for a life

without it, no matter how good the non-eudaimonic life may be in other respects. Derek Parfit nicely illustrates the spirit of this idea in the following passage:

I could live for another 100 years, all of an extremely high quality. Call this the Century of Ecstasy. I could instead live forever, with a life that would always be barely worth living [...] the only good things would be muzak and potatoes. Call this the Drab Eternity. I claim that, though each day of the Drab Eternity would be worth living, the Century of Ecstasy would give me a better life. Though each day of the Drab Eternity would have some value for me, no amount of this value could be as good for me as the Century of Ecstasy.

[Parfit 1984: 17-18]

Now suppose it can be shown that a eudaimonic life requires a certain amount of knowledge and that no amount of true belief that falls short of knowledge will do the trick. In that case, knowledge will also be weakly superior to true belief that isn't knowledge. After all, there will be an amount of knowledge that cannot be sacrificed for any amount of true belief that falls short of knowledge without losing something of superior value, to wit, the eudaimonic life.

To see how it works, let's first take a closer look at the eudaimonic life. Human beings are multi-faceted creatures leading complex intellectual, emotional and physical lives. Humans can attain flourishing in each of these domains. We believe that to attain a life of eudaimonia, one must attain flourishing in all of them. A human being leading a flourishing physical and intellectual life may fall short of a eudaimonic life if s/he does not attain a flourishing emotional life and so on.¹¹

We want to suggest that one important condition on flourishing in a given domain consists in sufficiently realizing one's potential in that domain. There are of course many ways in which one can fail to sufficiently realize one's potential in a given domain. Perhaps the most obvious one is to systematically and dramatically fail to attain what one should and could very easily have attained. Consider staying in reasonable physical shape. It is plausible that this is a goal

¹¹ Note that this is not to say that attaining flourishing life in a given domain involves an exceptionally high degree of flourishing in that domain. It also doesn't mean that one has to attain equal degrees of flourishing in all domains or that a eudaimonic life will be out of the question for anyone who attains significantly different degrees of flourishing in some of these domains.

that human agents should attain. Now consider an agent who has been gifted with a naturally fit physique. It is very easy for him to stay in shape. All he needs to do is keep to a minimally healthy diet and work out twice a week for a mere 15 minutes. However, he systematically and dramatically fails on this front with the result that he ends up pathologically obese and has to suffer all the physical impediments that accompany pathological obesity. There can be no question that this agent does not attain a life of physical flourishing. Moreover, given that physical flourishing is required for the eudaimonic life, this agent will also fall short of eudaimonia.

The same goes for the intellectual domain. That is to say, in order to attain intellectual flourishing one must sufficiently realize one's potential in that domain. Again, one way of failing to sufficiently realize one's potential in this domain is to systematically and dramatically fail to attain what one should and could very easily have attained. Now recall that knowledge is widely and easily available to us human beings in a wide range of areas. As a result, in these areas we humans often have the ability to know, knowledge is within our reach as intellectual agents. What's more, it is also plausible that among the things that it is easy for us to know, there is a wide range of things that we should know. If so, any human being who systematically and dramatically fails to attain knowledge that he could easily have attained will systematically and dramatically fail to attain knowledge of things in that range. As a result, any such human will fall short of sufficiently realizing his potential in the intellectual domain and so will fail to attain a life of intellectual flourishing. Given that a life of eudaimonia requires flourishing in all of the above mentioned domains, including the intellectual, any human agent who systematically and dramatically fails to attain knowledge s/he could easily have attained will be precluded from a life of eudaimonia. For us humans, then, the eudaimonic life is a life rich (enough) in knowledge.

6 Conclusion

In this paper, we have argued that knowledge is more valuable than any proper subset of its constituents because, in a wide range of areas, knowledge is widely and readily available. Knowledge is our way of correctly representing the world. In this way, we have addressed two important challenges that any satisfactory account of the value of knowledge will have to come to terms with, to wit, C1 and

C2. While we believe that Pritchard's additional challenge, C3*, is too strong, we granted Pritchard that the motivations for C3* he adduced might still serve to justify a demand on satisfactory accounts of the value of knowledge over and above C1 and C2, to wit, that knowledge must be taken off a value continuum with true belief that falls short of knowledge (C3). We provided some reason to think that C3 can be met as well by sketching an argument that life insufficiently rich in knowledge will not be a eudaimonic life.

References

- Goldman, A., & Olsson, E. 2009. Reliabilism and the value of knowledge. *In: Haddock, A., Millar, A., & Pritchard, D. (eds), Epistemic Value*. Oxford: Oxford University Press.
- Hazlett, A. 2010. The myth of factive verbs. *Philosophy and Phenomenological Research*, **80**, 497–522.
- Hursthouse, R. 2007. Virtue ethics. *In: Zalta, E.N. (ed), The Stanford Encyclopedia of Philosophy*, July 2007 edn.
- Kagan, S. 1998. Rethinking intrinsic value. *Journal of Ethics*, **2**, 277–297.
- Korsgaard, C. 1983. Two distinctions in value. *Philosophical Review*, **92**(2), 169–195.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lewis, D. 1996. Elusive knowledge. *Australasian Journal of Philosophy*, **74**, 549–567.
- Mill, J.S. 1963. *Essays on Ethics, Religion, and Society*. The Collected Works of John Stuart Mill, vol. 10. Toronto: University of Toronto Press.
- Parfit, D. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Pelletier, F.J. 2010. Mass Terms: A Philosophical Introduction. *In: Pelletier, F.J. (ed), Kinds, Things, and Stuff. Mass Terms and Generics*. Oxford: Oxford University Press.

- Pritchard, D., Millar, A., & Haddock, A. 2010. *The Nature and Value of Knowledge*. Oxford: Oxford University Press.
- Rabinowicz, W., & Rønnow-Rasmussen, T. 2000. A distinction in value: Intrinsic and for its own sake. *Proceedings of the Aristotelian Society*, **100**, 31–53.
- Radford, C. 1966. Knowledge–by examples. *Analysis*, **27**, 1–11.
- Riggs, W. 2009. Understanding, knowledge, and the Meno requirement. In: Haddock, A., Millar A. Pritchard D. (ed), *Epistemic Value*. Oxford: Oxford University Press.
- Rønnow-Rasmussen, T., & Zimmermann, M. (eds). 2005. *Recent Work on Intrinsic Value*. Dordrecht: Springer.
- Sosa, E. 2003. The place of truth in epistemology. In: DePaul, Michael, & Zagzebski, Linda (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*. OUP.
- Treanor, N. 2013. The measure of knowledge. *Noûs*, **47**, 577–601.
- Treanor, N. 2014. Trivial truths and the aim of inquiry. *Philosophy and Phenomenological Research*, **89**, 552–59.
- Zimmermann, M. 2015. Intrinsic vs. extrinsic value. In: Zalta, E.N. (ed), *The Stanford Encyclopedia of Philosophy*, spring 2015 edn.