

The Tertiary Value Problem and the Superiority of Knowledge

(DRAFT)

Mona Simion and Christoph Kelp

Abstract

In a number of recent pieces Duncan Pritchard has argued (TVP =) that any satisfactory account of the value of knowledge will have to solve the tertiary value problem; that is, it will have to show that knowledge is more valuable than true belief that falls short of knowledge as a matter of kind and not just as a matter of degree. One of the most promising candidates in the literature is John Greco's virtue epistemological solution according to which knowledge, as opposed to true belief, is a form of achievement and so, unlike true belief, finally valuable. In this paper, we will take a closer look at this debate over the tertiary value problem. We argue, first, that Greco's solution to the tertiary value problem remains unsatisfactory as he fails to show that true belief that falls short of knowledge is not finally valuable. Second, we will show that the motivation for TVP Pritchard offers do not actually support TVP. At best, they serve to establish a weaker demand on satisfactory accounts of the value of knowledge, one that can be met by showing that knowledge is weakly superior to true belief that falls short of knowledge. Finally, third, we offer two arguments for the weak superiority of knowledge. Since one of them rests on a Greco-style virtue epistemology, it also so serves to show that Greco can retain the upper hand in this debate.

1. Introduction

Knowledge has always been an object of special concern. This is evidenced by the amount of time and energy we invest in the development of institutions whose aim it is to accumulate or distribute knowledge. Also, in philosophy, a lot of effort has been made to get clear on what exactly is involved in knowing.

This gives rise to the following question: what is it about knowledge that is so distinctively valuable as to justify our special concern with it?

It is widely agreed that, in order to adequately answer this question, it must be shown that knowledge is more valuable than any lesser epistemic standing. But will that be enough? Suppose knowledge turned out to be only marginally more valuable than lesser epistemic standings. Would that vindicate our special concern with knowledge? As Duncan Pritchard has rightly pointed out, the answer to this question is plausibly ‘no’. That means that, in order to give an adequate account of the value of knowledge, we will have to do more. But how much more exactly? Pritchard’s own suggestion is that we will have to show that knowledge is more valuable than belief that falls short of knowledge not just as a matter of degree, but as a matter of kind.

Let’s suppose Pritchard is right: what we are looking for is an account of the value of knowledge on which knowledge enjoys a different kind of value than lesser epistemic standings. Enter John Greco. He argues that virtue epistemology (VE) can give us exactly what we need. According to a core thesis of VE, knowledge, as opposed to lucky true belief, is a cognitive achievement. Since

achievements are finally valuable, whilst lucky successes aren't, knowledge turns out to enjoy a kind of value not possessed by lesser epistemic standings.

This paper addresses the question as to whether Greco's account of the value of knowledge rises to Pritchard's challenge. Our ultimate aim is to defend a positive answer to this question in the sense that a virtue theoretic account of knowledge does serve to meet Pritchard's challenge properly understood. At the same time, the details of both the challenge and the virtue theoretic account will turn out to be rather different than Pritchard and Greco would have us think.

To this effect, we first outline Greco's account of the value of knowledge and how it serves to respond to Pritchard's challenge (§2). We will then turn to a recent objection to Greco's account due to Christian Piller, according to which achievements do not always have positive value. We discuss two lines of response to this objection, one unsuccessful, the other more promising (§3). While Greco can survive Piller's attack, we argue that his account faces another problem, to wit, it fails to establish at least one crucial claim, viz. that true belief is not of final value. As such, even if finally valuable, knowledge need not enjoy a different kind of value than lucky true belief (§4). Finally (§5), we argue that Pritchard's 'difference in kind' requirement is too strong and not supported by Pritchard's motivations, and that a virtue theoretic account of knowledge can rise to Pritchard's challenge properly understood.

2. Greco's Account of the Distinctive Value of Knowledge

The problem of explaining the distinctive value of knowledge dates back to Plato's *Meno* (1956); here, Socrates challenges his interlocutor to explain what

makes knowledge more valuable than mere true belief. After all, a true belief regarding the right way to Larissa is of just as much practical use as knowledge of the way to Larissa: both will get us where we want to go.

In recent years, several authors have rediscovered and generalized this challenge (also known as *the primary value problem*, henceforth C1). Jonathan Kvanvig (2003), for one, argues that if knowledge is analysable into simpler constituents, in order to account for our special concern with it, one must explain why knowledge is more valuable than any proper subset of its constituents, including, but not restricted to true belief (*the secondary value problem*, henceforth C2).

One could respond to the secondary value problem by arguing that knowledge is more valuable as a matter of degree than that which falls short of knowledge. It is unclear, however, whether this way of thinking about the value of knowledge can explain our special concern with it. Thus, Duncan Pritchard (2010) ups the stakes even further; according to him, in order to properly account for the distinctive value of knowledge, one must explain why knowledge is more valuable than any lesser epistemic standing not just as a matter of degree, but as a matter of kind. The reason for this is that if the difference in value is only a matter of degree, this has the effect of putting knowledge on a kind of continuum of epistemic value, albeit further up the continuum than anything that falls short of knowledge. The problem with this, however, is that it fails to explain why the long history of epistemological discussion has focused specifically on this particular point on the continuum. This has become known in the literature as *the tertiary value problem* (C3).

Notice that, if a response to the value problem turns out to be successful in meeting C3, it will thereby account for the primary and the secondary value problems too. As such, at least according to the desiderata one can identify in contemporary literature, by meeting C3, one will have given a fully satisfactory account of the distinctive value of knowledge.

One of the most prominent attempts to meet C3 comes from virtue epistemology (VE). The core thesis of this family of views is that knowledge is a cognitive achievement, a cognitive success that is due to cognitive competence or ability¹. Greco's core argument that virtue epistemology - on this construal - is able to explain the distinctive value of knowledge can be found in the following passages:

[An] answer to the value problem falls out of this account straightaway. In the *Nicomachean Ethics*, Aristotle makes a distinction between (a) achieving some end by luck or accident, and (b) achieving the end through the exercise of one's abilities (or virtues). It is only the latter kind of action, Aristotle argues, that is both intrinsically² valuable and constitutive of human flourishing. (2010: 97-98).

¹ For more on virtue epistemological accounts of the nature of knowledge see e.g. Sosa (2011), Greco (2010), Riggs (2002), Zagzebski (1996). One of us has also defended versions of virtue epistemology (e.g. Author 2011, 2013).

² Notice that Greco sometimes uses intrinsic and final value interchangeably (as does Aristotle). Crucially, however, contemporary value theory acknowledges that intrinsic value is not the same as final value. This is because intrinsic value concerns only the value generated by the intrinsic properties of the target item, and yet something can be finally valuable because of its relational (and hence non-intrinsic) properties. The standard example is Princess Diana's dress. In spite of having the same internal characteristics, it seems to be 1) more valuable than an exact replica, and 2) as opposed to the replica, valuable for its own sake. As such, the most recent approaches in value theory, i.e., 'Fitting Attitude' analyses of value, allow for final value to get inherited relationally (e.g. Anderson 1993, Rabinowicz and Ronnow Rasmussen 2004).

In several passages it becomes clear that what Greco means is actually final value: "I follow Aristotle in holding that success from ability is constitutive of human flourishing, which has final

And

[...] both success from ability and human flourishing have ‘final’ value, or value as ends in themselves, independently of any instrumental value they might also have. Therefore, knowledge has value over and above the practical value of true belief” (2010: 174f).

According to VE, then, knowledge is merely the cognitive incarnation of a more general, finally valuable kind: that of achievement. Thus, cognitive achievements will enjoy final value, as opposed to cognitive successes that are not from ability. If that is the case, one seems to be well on one's way to answering the tertiary value problem: knowledge enjoys value that is different in kind to that which falls short of knowledge.

With these points in play, it is not hard to see that the following three theses are key to Greco’s account of the value of knowledge:

The Knowledge = Achievement Thesis (KAT). Knowledge, as opposed to belief that falls short of knowledge, is a form of achievement

The Final Value of Achievement Thesis (FVAT). Achievements are finally valuable

The Exclusivity Thesis (ET). Successes that are not from ability fail to enjoy final value

value” (2010: 180). As such, throughout this paper, intrinsic value will be taken to refer to final value.

Several objections to Greco in the literature mainly focus on calling *KAT* into question; that is, the claim that knowledge is a form of achievement. Duncan Pritchard (2010) and Jonathan Kvanvig (2003), for instance, do not object to what may be called the value-theoretic part of an account along Greco's lines. Rather, they try to argue that the account fails because VE does not provide a correct account of the nature of knowledge. While arguments like Kvanvig's and Pritchard's have sparked a lively discussion on the prospects of virtue epistemological accounts of the nature of knowledge, it is fair to say that this debate is ongoing (see e.g. Greco 2010, Sosa 2011, Turri 2011, Author 2011 for rejoinders). In this paper, we will not take issue with the arguments on either side. Instead, we will grant Greco that there exists some workable virtue theoretic account of knowledge and focus on *FVAT* and *ET* in Greco's account.

3. The Final Value of Achievements Thesis

FVAT has been attacked by Christian Piller (2012) who argues that not all achievements are valuable: there are also worthless and bad achievements. Here is Piller: "[...] I cannot believe that there is anyone who would really think that every success from ability has something good in it. Think of the achievement of what is distasteful, brutal or, in some other way, horrible" (2012: 220). Take the case of Sid, a very skilful sniper, who is targeting five-year-old Charlie at a playground from a distance of 1000m. Despite the distance, his shot, let us

assume, hits Charlie right between the eyes. According to Piller, Sid produces “a good shot and a wicked and horrible act with nothing good in it” (2012: 220).³

3.1 Response 1: Overriding

Greco does not address this objection. However, one line of response to Piller that Greco might want to avail himself of can be found in Pritchard (2010); he argues that, properly understood, the thesis that achievements are finally valuable only claims that achievements have final value qua achievements, independent of the value of the outcome:

[...] some achievements may have no practical value, and may even accrue disvalue, perhaps because of the opportunity cost incurred by seeking the pointless achievement over a more substantive achievement or because of the wicked nature of the achievement in question. Indeed, there might be situations in which the all-things-considered value of Archie’s success of hitting the target when it is due to luck is much greater than the all-things-considered value of a corresponding success attained because of Archie’s ability. It is important to recognize that the value of achievements thesis when properly understood is entirely consistent with this possibility (2010: 32).

Let us attempt to unpack Pritchard’s claim. The core idea is that there are two types of value at work here:

³ One possible reply to this objection would be that, intuitively, we do not seem to call brutal and distasteful acts achievements. For instance, it sounds rather weird to refer to Jane’s cheating on her husband as an achievement, even though, for sure, Jane’s abilities might have been involved in her successfully doing so. Notice, though, that on Greco’s own account of achievement in terms of success from ability, this reply to Piller’s objection will not stand; successes from ability can surely come in both worthless and bad incarnations.

1. the final value of the achievement in question itself, qua achievement, which is taken to be (i) constant, (ii) positive and (iii) independent of any other type of value that may accrue to the resulting state of affairs
2. other types of value that may accrue to the resulting state of affairs (Pritchard himself focuses on practical value here).

With this distinction in play, it is easy to see that the all-things-considered value of a state of affairs involving an achievement may be null or even negative, even though the achievement makes some positive contribution to it. This will be the case when the positive value of the achievement qua achievement is outweighed by the negative value accrued by the resulting state of affairs. For instance, in the case of Sid and Charlie, the thought is that the added positive value of Sid's achievement qua achievement simply does not counterbalance the negative value of the loss of Charlie's life. In this way, *FVAT* does not preclude the all-things-considered value of the generated state of affairs being zero or even negative.

It may be worth noting that this goes even for cases in which the success from ability produces a good outcome. Suppose after two months searching, I finally find my missing green sock. Suppose, furthermore, that this is an achievement of mine: the search effort was conducted in a highly strategic manner, and was successful because of that. In this case, and unlike in the case of Sid and Charlie, the result I attain is good. Even so, the overall value of the state of affairs may still be negative: finding the sock was simply not worth the effort.

Finally, *FVAT* is also compatible with the overall value of a state of affairs involving a lucky success being higher than the overall value of the corresponding state of affairs involving a success from ability. For instance, in

the sock case, it would have been better if I had found the sock by sheer luck on the first day.

Crucially, however, it is far from clear that, in cases of wicked achievements, the achievement contributes positively to the overall value of the state of affairs, as Pritchard would have it. Thus compare the case of Sid, who achieves killing Charlie (Case 1), with a case in which Sid is also targeting him but, for some reason, produces a terrible shot, which misses Charlie by a long way. As it so happens, however, the bullet ricochets off the wall and hits Charlie right between the eyes anyway (Case 2). Notice that, according to Pritchard, the resulting state of affairs in Case 1 is better than in Case 2. To see this, notice that the cases are exactly parallel, except for the fact that, in Case 1, Sid attains achievement, whilst in Case 2 he only attains lucky success. Since achievements are supposed to contribute positively to the overall value, the value of the resulting state of affairs should be higher in Case 1. However, that does not appear to be the right result. Rather, if the resulting state of affairs isn't worse in Case 1, the two resulting states of affairs are at the very least equally bad in the two cases. Crucially, that means that Pritchard's attempt to rescue *FVAT* fails: in cases of wicked acts, the achievement is no better than the corresponding lucky success.

3.2 Response 2: Swamping

We think that there might be a more promising way for Greco to respond to Piller's worry. In a nutshell, the idea is that in cases of wicked achievements, the final value of achievement is not overridden by other aspects of the situation, but swamped by the wickedness of the achievement.

In order to get there, let's first look at another axiological property of achievements. Achievements come in degrees of quality: they can be of higher and lower levels. On the face of it, higher-level achievements are more highly valued than lower-level ones. Moreover, given that all else is equal between two achievements of the same kind, the level of difficulty determines the degree of quality of the achievement. If all else is equal between two successful shots in a shooting competition, the one that was more difficult to make is the higher-level achievement. Since the higher-level achievement is more valuable, the resulting state of affairs involving the achievement that was more difficult to attain accrues more value than the one featuring the less difficult achievement. Crucially, again, this does not hold for wicked achievements, at least not always. Suppose Sid's shot had been much more difficult, say, because he was shooting from a distance of 5000m and Charlie was moving around a lot (Case 3). In that case, Sid's kill constitutes a higher level of achievement than in Case 1. This, however, does not make things any better. In other words, the resulting state of affairs does not accrue more value in Case 3 than in Case 1.

It comes to light that wicked achievements differ from non-wicked ones in at least two respects. In the non-wicked case, (i) all else equal, the resulting state of affairs is better if it features an achievement than if it features the corresponding lucky success; furthermore, (ii) all else equal, it is better if it contains a higher-level achievement than a lower-level one. In contrast, in the wicked case, (i) all else equal, the resulting state of affairs is no better if it contains an achievement than if it contains the corresponding lucky success; and (ii) all else equal, it is no better if it features a higher-level achievement than a lower-level one.

This indicates that the axiological properties of achievement differ profoundly between non-wicked and wicked cases. How exactly do they differ? To answer this question, we'd now like to consider another case of final value: pleasure. To begin with, let's look at the non-wicked case. If you do something good, for instance, the resulting state of affairs will accrue more value if you derive pleasure from your doing good than if you don't, at least given that all else is equal. Moreover, all else equal, the resulting state of affairs will accrue more value the more pleasure you derive from it.

Let's move on to the wicked case. Suppose Kim brutally murders her innocent victim and derives an immense amount of pleasure from so doing (Case A). Contrast this with a case in which Kim also brutally murders her innocent victim but does not experience pleasure, say because she was given a drug that prevents her from experiencing pleasure (Case B). It is clear that the resulting state of affairs is no better in Case A than in Case B. Finally, consider a case in which Kim again brutally murders her innocent victim but derives only a lesser amount of pleasure from it, say because she was given a drug that decreases her capacity for experiencing pleasure (Case C). Again, there can be no question that the resulting state of affairs here is no worse than in Case A.

These considerations suggest that pleasure behaves (axiologically) just like achievement. Just as in the case of achievement earlier, in the non-wicked case, (i) all else equal, the resulting state of affairs is better if it features pleasure than if it doesn't; and, (ii) all else equal, it is better the more pleasure it contains. Moreover, in the wicked case, (i) all else equal, the resulting state of affairs is no better if it contains pleasure than if it doesn't; and, (ii) all else equal, it is no better if it contains more pleasure than less.

How do we explain this behaviour? As a first observation, notice that we will want to avoid the conclusion that pleasure is not finally valuable. After all, pleasure is perhaps the paradigm case of something that has final value. If not even pleasure is finally valuable, it is not clear any more that anything is. This means that we need an alternative explanation. We'd like to suggest the following hypothesis:

The Swamping Hypothesis (SH). If a state of affairs contains a wicked outcome that has been attained in a way that accrues final value, the final value accrued by the way is swamped by the wickedness of the outcome and so does not contribute to the total value of the state of affairs.

SH can accommodate the data in the wicked cases of both achievement and pleasure. After all, it predicts that, in cases of wicked outcomes that are attained by means of achievement/in a way that generates pleasure in the perpetrator, the resulting state of affairs is no more valuable than the same outcome attained via lucky success/without generating pleasure. Moreover, it also predicts that wicked outcomes attained via higher-level achievements/in a way that generates more pleasure is no more valuable than wicked outcomes attained via lower-level achievements/in a way that generates less pleasure.

At the same time, *SH* will allow Greco to hold on to *FVAT*. Like pleasure, achievements are finally valuable. It is just that, in cases of wicked achievement, they do not contribute value to the resulting state of affairs because the wickedness of the achievement swamps the value it would otherwise have

contributed. *SH* thus promises to give Greco a better way of responding to Piller than Pritchard did.

3.3 Some Consequences

There are a couple of further consequences of this way of responding to Piller that we'd like to point out. First, on the present response to Piller, it won't be the case that every item of knowledge is better than the corresponding achievement.

To see this, notice that there is very plausibly such a thing as wicked true belief. Consider true beliefs about how to most efficiently and brutally organise genocide and how to generate the most pleasure from torturing innocent infants. There is little room for doubt that these true beliefs are genuinely wicked. By *SH*, however, it follows that the resulting state of affairs will be no better if these true beliefs also constitute achievements (and hence qualify as knowledge) than if they are attained by luck.

One might think that this is bad news for Greco. After all, on the standard interpretation of the value problem, the challenge is to show that every item of knowledge is better than the corresponding true belief (Kvanvig 2003, Riggs 2009). If Greco accepts that there are some states of affairs involving knowledge that are no more valuable than the states of affairs featuring the corresponding lucky true beliefs, it may look as though Greco's account of the value of knowledge is still bound to fall short.

On the other hand, we think that this objection can now be plausibly turned on its head. After all, first, it is intuitively plausible that in cases of wicked true belief, the resulting state of affairs is no better if the true belief is achieved than if it is attained by sheer luck. Compare a case in which Bob achieves a true belief

about how to generate the most pleasure from torturing innocent infants (Case I) with one in which he acquires the corresponding gettiered belief (Case II). Intuitively, the resulting state of affairs accrues no more value in Case I than in Case II. Second, let's return to the case of pleasure once more. We take it that no one would be tempted to think that the fact that the final value of pleasure is swamped in cases of wicked pleasures indicates that pleasure is not distinctively valuable. In view of these considerations, then, it is more plausible that Riggs's way of unpacking the challenge is unduly demanding than that Greco fails to do it justice. The lesson to be learned here is that accounting for the distinctive value of knowledge does not require every item of knowledge to be more valuable than the corresponding lucky true belief.

Second, the present response to Piller's objection serves to confirm *FVAT*. After all, as we have just seen, achievement behaves axiologically like a paradigm case of something of final value, i.e. pleasure. Moreover, it behaves unlike a paradigm case of something of instrumental value, to wit, that of a Swiss army knife. A Victorinox Bentam is of lower instrumental value than a Victorinox Spartan as the latter possesses a number of tools the former doesn't. That said, when, during an adventure holiday, I slice my bread using a Bentam the resulting state of affairs is no worse than if I had done so with a Spartan. Unlike in the case of final value, the additional instrumental value of the Bentam is swamped even in the non-wicked case.⁴

4. The Exclusivity Thesis

⁴ This result is confirmed by other cases including, perhaps most notably, Linda Zagzebski's (2003) case of the cup of tasty espresso produced by a reliable vs. an unreliable espresso maker.

Let's move on to *ET*, the third crucial claim in Greco's argument. To begin with, we'd like to ask whether we always value positive achievements more than the corresponding lucky successes. The answer to this question appears to be 'no'. First, there are situations in which we seem to value success from luck at least as much as the corresponding achievement. For instance, it is not obvious that we value a student who has to work hard in order to get a certain set of skills more than one who is, luckily, naturally endowed with them. But couldn't this difference in value be explained in terms of the opportunity costs that attach to the achievement in comparison to the lucky success here? Again, we think that the answer here is 'no'. Even though the difference in opportunity cost clearly exist, it does not appear to fully capture the way in which we value the student's natural talent. Rather, we seem to attach special value to (luckily acquired!) natural talents in general.

Furthermore, there seem to be cases when we clearly value the lucky success more. Beauty is a case in point; as things stand at the moment, most beautiful people just luck on to their beauty. At the same time, it is possible to achieve beauty: one may repeatedly undergo plastic surgery. Crucially, in this case, we take it that nobody in their right mind would attach *more* value to achieving beauty than to lucking on to it. And, again, the difference in value does not seem to be explicable simply in terms of a difference in opportunity costs. It looks as though we attach special value to naturally acquired beauty. Finally, notice that this point generalizes to all incarnations of natural beauty: we seem to go through much trouble to preserve natural wilderness, for instance, when, in fact, we could just replicate it.

One move Greco could make at this stage would be to argue that, in all of the above cases, the extra value accrues relationally from the final value we attach to untouched nature in general. And, plausibly enough, he would be right. Notice, however, that this proves an important point: while tokens of positive achievement are taken to inherit final value from the finally valuable type, lucky success too can turn out to be finally valuable on independent grounds. That is, depending on the subject matter, lucky success can inherit final value relationally, independently of its not being an achievement. If so, *ET* is bound to be false as well.

In the light of this result, let us turn to the case of cognitive achievements now and assume that, in spite all difficulties, we grant Greco that knowledge is finally valuable. Does that necessarily mean that his account meets C3? That is, does that mean that we will have a difference in kind between knowledge and lesser epistemic standings? No. To see this, note that it is widely agreed that true belief is a fundamental value in the domain of epistemology, that is, a value that cannot be reduced to other distinctively epistemic values.⁵ If, as we think is plausible, the value of true belief is not reducible to fundamental values in other domains either, there is strong reason to believe that true belief is finally valuable.

Note that the situation here is, again, similar to the natural beauty examples, in that final value is accrued from two different sources: knowledge gets it in virtue of being the instantiation of a finally valuable kind – that of positive

⁵ In fact, many people think that true belief is the fundamental value in epistemology (Goldman 1999, Sosa 2007, David 2005, Olsson 2007, Ahlstrom-Vij 2013).

achievement – while lucky true belief inherits it relationally⁶ from the final value of truth.

Since, for present purposes at least, it is safe to assume that there is no difference in other kinds of value among knowledge and true belief, it now looks as though there is no kind of value that attaches to knowledge that does not attach to true belief. If so, C3 will of course not be met.

Greco might respond to this objection by arguing that, if truth is finally valuable, surely not only lucky true beliefs, but instances of true belief from ability too will inherit final value from this source. As such, cognitive achievements will turn out to have more final value attached to them than mere true beliefs, because they would inherit it from two sources⁷.

Unfortunately, this would only help Greco meet C1, and maybe C2⁸. That is, he would have proven that knowledge as cognitive success from ability has more final value attached to it than belief that falls short of knowledge. One thing is certain, though: by taking this route, Greco's account doesn't seem successful in meeting C3 anymore. Recall that Pritchard required that, in order to explain our special concern with knowledge, any satisfactory response to the value problem should account for knowledge enjoying value that is different in kind⁹ to the one attached to lesser epistemic standings. If both knowledge and true belief turn out

⁶ Recall that on the most recent approaches in value theory, final value gets inherited relationally. See note 2.

⁷ Greco himself uses the idea that some good A can have more final value than some good B when he discusses the value of understanding and wisdom vis-à-vis knowledge.

⁸ In what C2 is concerned, the question is left open. People might still argue that justification, for instance, has some further, independent final value attached to it too.

⁹ One way in which Greco might still hold on to the difference in kind idea is by arguing that the value that accrues to knowledge in virtue of its being an achievement is *sui generis*. That is, value of its own kind, which, as such, is different in kind to the final value of true belief. Notice, however, that it is not clear how this will still manage to meet Pritchard's worry regarding the motivation behind our special concern with knowledge. Thus, the question that naturally arises is what makes the kind of final value enjoyed by knowledge so distinctively more worthy of our concern than the corresponding kind enjoyed by belief that falls short of knowledge.

to be finally valuable, however, even if knowledge enjoys more final value, the envisaged difference in kind disappears.

5. Back to C3

We have just seen that, even if one manages to account for knowledge being finally valuable, that would not necessarily imply meeting C3. That is because lesser epistemic standings might turn out to be finally valuable too. By the same token, however, one might think that this merely serves to show is that C3 is too demanding. After all, if it can be shown that knowledge is finally valuable, shouldn't this be enough to adequately address any legitimate challenge concerning the value of knowledge?

While we sympathize with this line of thought, we also think that Pritchard may have been on to something when he introduced C3. To see this, let's take another 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 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.

5.1 Superiority in 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 belief that falls short of knowledge, it isn't necessary. There are other ways in which the value of one type of good, A, can be discontinuous with the value of another type of good, B. Mill famously put forth such a discontinuous account of value relations¹⁰:

It is quite compatible with the principle of utility to recognize 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

¹⁰ Many contemporary axiologists follow suit (see e.g. Parfit 1984, Rabinowicz and Arrhenius 2003).

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).

On a Millian reading, there are two ways in which the value of one type of good, A, can be discontinuous with the value of another type of good, B, without requiring a difference in kind: 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’). For illustration, consider how no finite amount of some low-quality value – having thousands of cups of good espresso, for instance – could compensate for losing some high-quality value as, say, the love of one’s child. As such, the value of the latter is strongly superior to the value of the former.

There are also cases where the discontinuity in value ordering is less abrupt. Here is an illustration by Derek Parfit:

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 (1984: 17-18).

Notice that Parfit's argument does not support Strong Superiority. In order to do that, Parfit would have to argue, for instance, that instead of living another 20 years from now on muzak and potatoes, he would rather die after an instant of ecstasy. And this does not sound very compelling. However, his case does illustrate Weak Superiority: some amount of high quality years is plausibly preferable to any amount of barely worthwhile ones.

Recall that, according to Pritchard, the only way in which we can avoid placing knowledge on a value continuum is by showing that it enjoys a different kind of value. By now, it has become transparent that this claim is false. Both Strong and Weak Superiority will take knowledge off a continuum with belief that isn't knowledge. At the same time, neither requires a difference in kind between these two. In order to account for the special value of knowledge in a way that does justice to Pritchard's motivation for C3, it should be sufficient to show that knowledge is, in some respect, or to some purpose, at least weakly superior to belief that falls short of knowledge. That is to say, it will do to show that some amount of knowledge is at least in some respect, or to some purpose, better than any amount of true belief that falls short of knowledge.

5.2 VE and the Weak Superiority of Knowledge

To see how this can be done, let us first return to Greco's argument for the thesis that achievements enjoy final value. Recall that Greco takes after Aristotle in distinguishing between the value of (a) achieving some end by luck or accident, and (b) achieving the end through the exercise of one's abilities or virtues. It is only the latter kind of action, Aristotle argues, that is both intrinsically (i.e., finally) valuable and constitutive of human flourishing.

Let us try to disambiguate this a bit. Notice that one can distinguish two separate claims by Aristotle here:

The Exclusivity of Final Value Thesis (EFVT). Only achievements, and not lucky successes, are intrinsically (i.e., finally) valuable.

The Exclusivity of Human Flourishing Thesis (EHFT). Only achievements, and not lucky successes, are constitutive of human flourishing.

In the light of our results in the previous sections, it is now easy to see that *EFVT* will not stand. After all, we found that there are cases in which lucky successes can be finally valuable too, in virtue of their standing in a relationship to something else we find valuable. Recall, for instance, how the final value of unaltered nature got inherited in cases of (luckily acquired) natural beauty or talents. As such, it seems that it is not the case that only achievements and not lucky successes are finally valuable. Thus, *EFVT* is too strong.

A similar fate befalls *EHFT*. As cases of natural beauty plausibly indicate, lucky successes may also be (at least partly) constitutive of an individual's flourishing. What remains plausible, however, is the weaker:

The Necessity of Achievement Thesis (NAT). A life insufficiently rich in achievement will not be a life of human flourishing.¹¹

¹¹ We do not mean to suggest that the critical threshold of achievement here is simply a quantitative matter. It would seem that one can achieve many trivial or bad things without making much headway in the direction of human flourishing. In what follows, we will take any further restriction on the kinds of achievement required for human flourishing (good, significant, etc.) as read.

To see the plausibility of *NAT*, consider Lucy who has had the good fortune of being gifted with a lot of natural endowments, including beauty and a wide range of natural talents. In addition, she was born into family that is wealthy enough that Lucy will be able to live very comfortably for her entire life. Lucy is one lucky woman if ever there has been one. Now suppose that Lucy goes through life without ever achieving anything of her own. The things she does endeavour to achieve she either fails to achieve or attains by luck. It is very plausible that Lucy fails to attain a life of human flourishing. Why? One plausible answer is that a life of human flourishing requires one to sufficiently realize one's potential and Lucy fails spectacularly on this front.¹²

In addition, the following is also plausible:

Weak Superiority of Human Flourishing (WSHF). A life of human flourishing is weakly superior to a life without it.

WSHF claims that there is some amount of human flourishing such that a life that has that amount is better than a life without it, no matter what other good-making properties may attach to this life. But now notice that, in conjunction with *WSHF*, *NAT* supports the thesis that achievement is weakly superior to lucky success in the sense that a life rich enough in achievement is weakly superior to a life with any amount of lucky success. After all, by *NAT*, exchanging the former

¹² Note that Pritchard also endorses *NAT*, or at least something very close to it, in the following passage: "Indeed, that achievements are valuable in this way is hardly surprising, once one reflects that they constitute the exercise of one's agency on the world. A life lacking in such agential power, even if otherwise successful (e.g. one's goals are regularly attained), would clearly be severely impoverished as a result. A good life is thus, amongst other things, a life rich in achievement" (2010: 30-31).

for the latter would mean jeopardizing a superior good, or at least any prospect thereof.

We human beings are multi-faceted. We have complex physical, emotional, cognitive, etc. lives. Moreover, we can attain flourishing in each of these domains.¹³ Let's focus on our cognitive lives and ask the question what is required for a flourishing cognitive life. The answer, we'd like to suggest, is parallel to the one we gave in the case of human flourishing:

The Necessity of Cognitive Achievement Thesis (NACT). A life insufficiently rich in cognitive achievement will not be a flourishing cognitive life.

In addition, the parallel weak superiority thesis is also plausible:

Weak Superiority of Cognitive Flourishing (WSCF). A flourishing cognitive life is weakly superior to a cognitive life without flourishing.

It is easy to see that we can now argue in a parallel fashion that, *for the purpose of leading a flourishing cognitive life*, cognitive achievement is weakly superior to lucky cognitive success, in the sense that a life rich in cognitive achievement is weakly superior to a life with any amount of lucky cognitive success. Given Greco's virtue epistemology, which identifies knowledge with cognitive achievement (*KAT*) and lucky success with true belief that falls short of

¹³ Is it possible to attain flourishing in one domain without attaining flourishing in all of them? Very plausibly so. Is it possible to reach the critical threshold in a domain required to attain superior human flourishing without attaining the critical threshold required for human flourishing in all domains? While this is less clear, it is still plausible. Can one attain human flourishing if one has not attained the critical threshold for flourishing across all domains? This is perhaps least evident. For present purposes, we will leave these questions open.

knowledge, we get the result that, *for the purpose of leading a flourishing cognitive life*¹⁴, knowledge is weakly superior to true belief that falls short of knowledge.

5.3 Another Route to the Weak Superiority of Knowledge

The above argument for the weak superiority of knowledge over true belief that falls short of knowledge relies on virtue epistemology, and, in particular, versions of the view that endorse *KAT*. It may be worth noting, however, that there is another argument that will work on any account of knowledge. To see how it works, let's return to a fundamental condition on flourishing, to wit, that one sufficiently realizes one's potential. There are many ways in which one can fail to sufficiently realize one's potential in a given domain, the most obvious one of which is systematically and dramatically failing to attain what one could very easily have attained. Consider Heathcliff 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 obese and has to suffer all the physical impediments that accompany obesity. There can be no question that Heathcliff does not attain a life of physical flourishing.

¹⁴ Notice that we are not claiming that some amount of knowledge is better than any amount of true beliefs in absolute terms. This claim would be obviously problematic; for practical purposes, for instance, having any amount of knowledge that is practically useless will not be more valuable than having some amount of practically useful true beliefs. Our claim is weaker: we hold that, given that a cognitively flourishing life requires some amount of cognitive achievements, *for the aim of leading a cognitively flourishing life*, some amount of knowledge is better than any amount of true beliefs.

Crucially, as regards the cognitive domain – at least for us, in the kinds of epistemic environments we inhabit – knowledge is regularly within reach and often very easily so.¹⁵ All I need to do to come to know that there is milk in the fridge is go and have a look. If I want to come to know whether my friend is coming to tonight’s concert, all I need to do is give him a call. If I want to know whether the meeting is at 4pm tomorrow, all I need to do is check my agenda. Cognitive agents who only ever form belief that falls short of knowledge will systematically fail to attain what they could very easily have attained. They will suffer from the very sort of systematic and dramatic failure that is incompatible with a life cognitive flourishing. For us, a life of cognitive flourishing is a life rich (enough) in knowledge.

All that we need to complete the argument is the point, made earlier already, that a life of cognitive flourishing is weakly superior to a life without it. This will give us the result that knowledge is weakly superior to belief that falls short of knowledge (again, for the purpose of leading a life of cognitive flourishing), in the sense that a life rich enough in knowledge is weakly superior to a life with any amount of belief that falls short of knowledge, at least for us, in the kinds of epistemic environments we inhabit.

6. Conclusion

It comes to light that the most promising response to the tertiary value problem, coming from John Greco, fails to serve its purpose in the way envisaged by its proponent. Greco purports to meet C3 by showing that there is a difference in

¹⁵ We argue elsewhere (Author and Author 2014) that, in most areas, knowledge is more readily available than belief that falls short of knowledge.

kind between knowledge and lesser epistemic standings. Greco analyses knowledge, as opposed to lucky success, in terms of cognitive achievement (*KAT*), and follows Aristotle in arguing that achievements enjoy final value (*FVAT*), whereas lucky successes don't (*ET*). We have argued that Greco's account of the value of knowledge carries quite a bit of promise. *KAT* is still a live option in the debate over the nature of knowledge. Moreover, we have argued that there is a promising response to a major objection against *FVAT* available to Greco. Even so, *ET* turned out to remain problematic. As stated, then, Greco's account does not succeed in meeting C3.

The main ambition of this paper, however, was to show how this objection is not fatal to Greco's account. We have argued that, properly understood, the motivations for C3 do not justify the requirement of showing that knowledge has a different kind of value. The kind of discontinuity in value ordering at issue there will also be obtained if knowledge turns out to be at least weakly superior to true belief that falls short of knowledge. We have argued that Greco's *KAT* can still be used to argue that knowledge is weakly superior to true belief that falls short of knowledge. In this way, *KAT* contributes to taking knowledge off a value continuum with lesser epistemic standings, without thereby accounting for a difference in kind of value between the two.

References

Author (2011)

Author (2013)

Author and Author (2014)

- Ahlstrom-Vij, K. (2013). In Defense of Veritistic Value Monism. *Pacific Philosophical Quarterly*, 94, 19-40.
- Anderson, E. (1993). *Value in Ethics and Economics*. Cambridge, Mass: Harvard University Press.
- David, M. (2005). Truth as the Primary Epistemic Goal: A Working Hypothesis. In E. S. Steup (Ed.), *Contemporary Debates in Epistemology* (pp. 296–312). Oxford: Blackwell.
- Goldman, A. (1999). Knowledge in a Social World. *Argumentation* 16, 3, 369-382.
- Greco, J. (2009). The Value Problem. In A. M. A. Haddock (Ed.), *Epistemic Value* (pp. 313-321). Oxford: Oxford University press.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Mill, J. S. (1963). Essays on Ethics, Religion and Society. In J. Robson (Ed.), *The Collected Works of John Stuart Mill, vol. 10*. Toronto: University of Toronto Press.
- Olsson, E. J. (2007). Reliabilism, Stability and the Value of Knowledge. *American Philosophical Quarterly* 44 (4), 343-355.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Piller, C. (2012). Knowledge as Achievement - Greco's Double Mistake. In W. L. C. Jaeger (Ed.), *Epistemology: Contexts, Values, Disagreement* (pp. 215-225). Mentis.

- Plato. (1956). The Meno. In W. Guthrie (Ed.), *Protagoras and Meno* (W. Guthrie, Trans.). London: Penguin.
- Pritchard, D.H., Millar, A. and Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press.
- Rabinowicz, W. a. (2003). On Millian Discontinuities. In W. R. Rønnow-Rasmussen (Ed.), *Patterns of Value: Essays on Formal Axiology and Value Analysis* (pp. 1-8). Lund: Lund Universitetstryckeriet.
- Rabinowicz, W. a.-R. (2004). The Strike of the Demon: On Fitting Pro-Attitudes and Value. *Ethics*, 114, 391–423.
- Riggs, W. (2002). Reliability and the Value of Knowledge. *Philosophy and Phenomenological Research* 64, 79-96.
- Riggs, W. (2009). Understanding, Knowledge and the Meno Requirement. In A. Haddock, A. Millar and D.H. Pritchard (eds.) *Epistemic Value* (pp. 331-38). Oxford: Oxford University Press.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
- Sosa, E. (2011). *Knowing Full Well*. Princeton: Princeton University Press.
- Turri, J. (2011). Manifest failure: The Gettier problem Solved. *Philosophers' Imprint*, 11, 1-11.
- Zagzebski, L. (1996). *Virtues of the Mind*. Cambridge: Cambridge University Press.
- Zagzebski, L. (2003). Intellectual Motivation and the Good of Truth. In M. DePaul and L. Zagzebski (Eds.), *Virtue: Perspectives from Ethics and Epistemology* (pp. 135-54). Oxford: Oxford University Press.

