Philos Stud
DOI 10.1007/s11098-010-9620-2

Reliabilism and the extra value of knowledge

Wayne A. Davis · Christoph Jäger

Abstract Goldman and Olsson (2009) have responded to the common charge that reliabilist theories of knowledge are incapable of accounting for the value knowledge has beyond mere true belief. We examine their "conditional probability solution" in detail, and show that it does not succeed. The conditional probability relation is too weak to support instrumental value, and the specific relation they describe is inessential to the value of knowledge. At best, they have described conditions in which knowledge indicates that additional epistemic value is likely to be forthcoming in the future. We also argue that their motive analogy breaks down. The problem, we conclude, is that being produced by a reliable process is not sufficient for a belief to be justified.

Keywords Knowledge · Extra value · Reliabilism · Justification · Truth · Swamping problem

1 The extra value of knowledge problem

Why is knowledge valuable? Part of the answer, plausibly, is that knowledge is at least true belief, which is valuable.[1] It is generally good for our beliefs to be true. But the cases which show that knowledge is more than true belief also make it plausible

[1] See Plato, Meno, 97a.

W. A. Davis (✉)
Georgetown University, Washington, DC, USA
e-mail: davisw@georgetown.edu

C. Jäger
University of Innsbruck, Innsbruck, Austria
e-mail: christoph.jaeger@uibk.ac.at

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

that knowledge has some extra value.[2] These are of at least three kinds: unjustified true beliefs; Gettier cases; and lottery cases. (i) Suppose a police officer infers that Carlos is a criminal from the fact that Carlos is a gang member and some gang members are criminals. Suppose further that Carlos is a criminal. Then the officer's belief is true, but it is not knowledge. It is not knowledge because the officer's reasoning is fallacious. As a result, her belief is unjustified. It was formed improperly. Given that Carlos is a criminal, it would be better for the officer to know that Carlos is one than for her to improperly believe that he is. For certain practical purposes (e.g., apprehending criminals), it may not matter whether the officer has knowledge or true belief.[3] But this difference does matter to our evaluation of the officer and her actions. (ii) Suppose that a juror infers that one of the defendants is guilty from the fact that a particular defendant confessed to the crime. The juror's belief is true because one of the defendants is guilty. However, the defendant who confessed is not the one who committed the crime. Then even though the juror has a true belief, he lacks knowledge. He is right for the wrong reason. Given that the juror's belief is true, it would be better for him to know that one of the defendants is guilty than to be right for the wrong reason. (iii) Steve believes that he will lose the lottery because the odds against winning are astronomical. As a result, he is tempted to sell the ticket. In fact, he will lose. Even though Steve's belief is true, he does not know that he will lose. For he knows he has a chance of winning. Consequently he has some doubt about the outcome and his evidence makes it reasonable for him to have that doubt. Given that Steve's belief is true, it would be better for him to know that he will lose than to have a reasonable doubt that he will. The Gettier and lottery problems are notoriously difficult, so we will focus primarily on the requirement that knowledge be justified true belief.

We do not need to assume that knowledge is always more valuable than mere true belief.[4] Our goal here is to determine whether process reliabilism can account for cases in which knowledge does have extra value.[5] The central idea of process reliabilism is that knowledge is true belief resulting from a reliable process. An additional clause may be added to exclude Gettier cases. Roughly speaking, a reliable process is one that produces mostly beliefs that are true (Goldman and

2FL01 [2] While this assumption is widely shared (in addition to Goldman and Olsson 2009, see Jones 1997;
2FL02 Zagzebski 2000, 2002, 2003; Riggs 2002; Sosa 2003; Brogaard 2006; Brady 2006; Olsson 2007;
2FL03 Pritchard 2007a, b; Pritchard et al. 2010, Chap. 1; Greco 2010, Chap. 6), it has been denied (e.g., by
2FL04 Swinburne 1999, 2001; Kvanvig 1998, 2003; Baehr 2009, unpublished). Our goal here is not to provide a
2FL05 full defense of the assumption but to see whether process reliabilism could account for it.

3FL01 [3] For a detailed discussion of the alleged importance of knowledge for practical reasoning, see Weiner
3FL02 2009. Weiner argues that beliefs that constitute knowledge need not be better premises for practical
3FL03 reasoning than merely true beliefs.

4FL01 [4] Baehr (2009b; see also Baehr 2009a) argues that when it comes to trivial subject matters (such as the
4FL02 number of blades of grass on your neighbor's lawn), knowledge is not epistemically more valuable than
4FL03 mere true belief. However, for those who think epistemic value does not reduce to some form of practical
4FL04 value, Baehr's claim is debatable. In any event, Baehr's argument leaves untouched the value problem as
4FL05 it arises for beliefs about non-trivial subject matters.

5FL01 [5] Zagzebski, e.g., in 2000, 2002, 2003; Riggs 2002; Sosa 2003; Greco 2002, 2003, 2010, Chap. 6; and
5FL02 others have argued that, contrary to process (and some other forms of) reliabilism, virtue reliabilism is in
5FL03 a position to solve the value problem. For a critical discussion of this view see, for example, Brogaard
5FL04 2006 and Lackey 2007.

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

Olsson 2009). The problem is that the property of being produced by a process that produces more true beliefs than false beliefs does not seem capable of adding any value to a true belief. A reliable process is valuable only because it produces beliefs that have the independently valuable property of corresponding to the facts. The production process does not give the beliefs that value. And a true belief produced by a reliable process is no better than a true belief formed in other ways.[6]

2 The motive analogy

Goldman & Olsson suggest an analogy designed to support the claim that being produced by a reliable belief-forming process adds value to a true belief. They observe that in addition to valuing actions, we also value motives. Doing something with a good motive is better than doing it with a bad motive. Goldman & Olsson ask why we value motives, and suggest that "a straightforward explanation is that such motives regularly bring about corresponding actions, actions which themselves are valuable." They observe further that:

> it is very plausible that good motives or intentions are *among* the things rated as *independently* good. This is confirmed by intuitive judgements to the effect that a compound state consisting of a good motive and a good action is (morally) *better* than a compound state consisting of the same good action done from a bad (or non-good) motive. Apparently, a good motive's value can be added to the value of a good action…. (Goldman and Olsson 2009, p. 33).

Suppose for concreteness that a surgeon saves a man's life by amputating his leg. Then amputating the patient's leg was a good thing to do. The action may have been done with different motives, however. If the surgeon's motive in amputating the leg was to save the patient's life, that is better than if her motive was to torture the patient. Goldman and Olsson suggest that the motive of saving lives is valuable only because it typically causes agents to save lives. This explanation of the motive's value does not seem adequate. Causing some good is not sufficient to make a motive good. If the motive of torture regularly resulted in agents saving lives, that would do little to make it a good motive. The motive of saving lives is good at least in part because having that motive contributes to having a benevolent attitude toward others, and thus contributes to being a good person. Causing some good is not essential for a motive to be good either. Given the value of human life, we should want to save lives even if we never have the opportunity to do so. Reliable belief-forming processes do not have independent value of this sort.[7]

6FL01 [6] For influential presentations of this "swamping problem," see the works quoted in footnote 2. The
6FL02 popular label is due to Kvanvig.

7FL01 [7] Goldman & Olsson might suggest that the motive is valuable for someone who never has the
7FL02 opportunity to save a life because it would likely cause some good if that opportunity arose. Given all the
7FL03 counterfactual circumstances in which that opportunity could arise, however, it is hard to see how that
7FL04 likelihood claim could be defended. Furthermore, if it were true, the most that would seem to follow on
7FL05 an instrumentalist view of value is that the motive *would be good* if those opportunities arose, not that it *is*
7FL06 good.

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

There are two further problems. First, the surgeon's motive may have been very specific: saving this patient's life here and now. The value of such a specific motive cannot be due to its regularly having good effects. Second, the action the surgeon performs with a good or bad motive is different from the object of the motive. The surgeon's motive for amputating a leg must be something other than amputating the leg. In our example, the surgeon expects to save the patient's life as a result of amputating the leg. So the fact that a good motive's value is added to the value of a good action does little to support the conclusion that the value of a true-belief producing process adds to the value of a true belief. The alleged analogy fails.

3 The validity of the conditional probability solution

Goldman and Olsson propose that the property that makes knowledge more valuable than true belief is:

> the property of *making it likely* that one's future beliefs of a similar kind will also be true. More precisely, under reliabilism, the probability of having more true belief (of a similar kind) in the future is greater conditional on S's *knowing* that p than conditional on S's *merely truly believing* that p (Goldman and Olsson 2009, p. 28).

There are a number of problems with this *conditional probability solution* (CPS). The first is the inference that a factor adds value from the fact that it raises the conditional probability of a good. Goldman and Olsson's inference can be represented as follows, displaying its general form:

CPS G and G′ have value.
$P(G'/G\&F) > P(G'/G\&\neg F)$
$\therefore$ G&F has more value than G.[8]

In Goldman & Olsson's intended application, G states that the subject has a particular true belief, which is an instance of knowledge. We want to know why its being knowledge makes it more valuable than a mere true belief. G′ states that the subject will have more true beliefs of a similar kind. F is:

> F G resulted from a reliable belief-forming process.

Let us assume for now that the probability premise of CPS is true. Does the extra value conclusion follow? It may seem like good means-ends reasoning, with the extra value being instrumental. But in fact, the conditional probability relation is too weak to support the conclusion. We can see this by considering alternatives to F, such as the following.

[8] More precisely, the conclusion is the proposition that the state of affairs described by G&F is more valuable than that described by G. F similarly is the proposition that the belief ascribed by G resulted from a reliable process. In order to avoid terminological complexity and excessive formality, we are using capital letters equivocally for both propositions and the states of affairs those propositions represent.

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

F_a The subject will successfully commit a series of bank robberies.
F_b G resulted from a reliable belief-forming process that sometimes malfunctions with fatal consequences.
F_c G′, and the subject will acquire an even greater number of false beliefs.
F_d G′, and any beliefs the subject acquires in addition to G will be unjustified.
F_e G resulted from a reliable belief-forming process that never produces justified beliefs.

Consider F_a, and let G state that the subject has the true belief that he is in a bank. Since it is impossible to successfully rob banks without forming a large number of true beliefs about banks and their contents, F_a makes G′ extremely probable. More specifically, the probability of G′ given G&F_a will be greater than the probability of G′ given G&$-F_a$ in a wide variety of circumstances if not all.[9] This conditional probability statement is at least as plausible as Goldman and Olsson's. Yet it does not follow that F_a is good, nor that G&F_a is more valuable than G alone. The problem in this case is that the added factor has a negative value that may far exceed the positive value of a future true belief.

Since having a reliable belief-forming process seems like a good thing, it may seem easy for Goldman and Olsson to strengthen the premises of CPS to exclude factors like F_a. But how can they block F_b? The reliabilist's paradigm cases of reliable belief-forming processes, such as perception and memory, sometimes do malfunction with fatal consequences. If Goldman and Olsson count the future benefits of the reliable processes when accounting for the extra value of the true beliefs they produce, how can they avoid counting the costs of relying on the processes?

The reliabilist might try to avoid these problems by replacing 'value' with '*epistemic* value.' It will be difficult, however, to define epistemic value before settling the question of why knowledge has more value than mere true belief. More importantly, CPS seems invalid even when the conclusion is restricted to epistemic value and epistemic value is tied tightly to truth. Consider F_c, which entails that the subject acquires a number of false beliefs greater than the number of true beliefs described by G′. This is compatible with the claim that the false beliefs are produced by the same process that produced G, since a reliable process can produce a finite run of bad results. Given that F_c entails G′, P(G′/G&F_c) has the maximum value, 1. P(G′/G&$-F_c$) will be less than 1 except in very extraordinary circumstances. In all cases, however, it seems that G&F_c has lower epistemic value than G.

Consider F_d. Given that F_d also entails G′, P(G′/G&F_d) will be greater than P(G′/G&$-F_d$) except in the most extraordinary circumstances.[10] Yet it is difficult to decide whether the conclusion of CPS in this case is ever true. It seems plausible that being unjustified is a negative epistemic value that can counteract or outweigh

[9] For the relation to fail, circumstances would need to be such that if the robbers do not bring off a series of bank robberies, then they are certain to do something else that requires true beliefs about banks. Such circumstances would obtain only if the subject is certain not to instantly drop dead of a heart attack or something else.

[10] Note that $-F_d$ is not the statement that the additional beliefs formed are true and justified. That statement is a contrary of F_d but not its negation.

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

163 truth. Is having additional beliefs that are true but unjustified more (epistemically)
164 valuable than having just one true belief?
165 In the case of F_e, it seems clear that the conclusion would be false despite the
166 probability premise being as plausibly true as Goldman and Olsson's. The reliabilist
167 may respond here that F_e describes an impossibility because beliefs that result from
168 reliable processes are justified (see e.g., Goldman 1979). However, it is easy to
169 describe processes that produce mostly true beliefs despite the fact that the subject
170 makes fallacious inferences, has undermining counterevidence, or fails to get the
171 evidence needed.[11] Goldman (1986, p. 110) himself has subsequently argued that
172 being produced by a reliable process is necessary but not sufficient for being
173 justified. What the possibility of F_e being true indicates is that the extra value of
174 knowledge is due to the belief's being justified. Its being produced by a reliable
175 belief-forming process is not sufficient. The reliabilist might try to insist that being
176 unjustified is not epistemically bad, but that is difficult if 'epistemic' means
177 anything like "of or pertaining to knowledge."
178 The most that would appear to follow from the premises of CPS is that G&F
179 *indicates* that additional value will be forthcoming in the future. To say that G&F is
180 an indicator of additional value is not to say that it has the additional value.
181 Goldman & Olsson's goal, however, is to explain why knowledge has more value
182 than mere true belief.

183 4 The soundness of the conditional probability solution

184 We have cast doubt on the validity of CPS. We now argue that its probability
185 premise is unwarranted. The assertion that $P(G'/G\&F) > P(G'/G\&{-}F)$ is based on
186 two important assumptions. Let *r* be the process G resulted from. The first
187 assumption is that if something good results from a process, then the subject is
188 likely to reuse that process. Specifically:

189 (1) If G results from process *r*, then the subject will reuse *r*.

190 The second assumption is that if the subject will reuse *r*, then G' is more likely if *r* is
191 reliable.

192 (2) Given that the subject will reuse *r*, G' is more likely if *r* is reliable than if it is not.

193 The second assumption is intuitively plausible, and seems to support inequality (3),
194 and thereby the probability premise of CPS. Let R be the proposition that the subject
195 will reuse *r*.

196 (3) $P(G'/R\&F) > P(G'/R\&{-}F)$.

11FL01 [11] A classic example is BonJour's clairvoyant, Norman, who always forms, in a reliable fashion, true
11FL02 beliefs about the current whereabouts of the US president, yet has no positive evidence for the fact that he
11FL03 possesses these special cognitive capacities. In other examples, the subject even has undermining
11FL04 counterevidence against the reliability of his (in fact reliable) cognitive processes and faculties and thus
11FL05 has counterevidence against the truth of what he or she comes to believe through these processes and
11FL06 faculties (see BonJour 1985, pp. 41 ff.).

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

However, −F is the counterfactual statement that G did not result from a reliable process. This does not entail that *r* is unreliable. For −F would also be true if *r* were reliable but did not produce G. Given that the subject will reuse *r*, the probability of G′ depends on the reliability of *r*, and not on whether it produced G. The truth of (2) is quite compatible with (4):

(4) $P(G'/R\&F) = P(G'/R\&{-}F) = P(G'/R)$.

But if (4) is true, then there is no basis for the claim that F raises the probability of G′. That is, there is no reason to assume that G′ is more likely given G&F than given G&−F.

To see a further problem, let us stipulate that F makes it likely both that the subject will reuse *r* and that *r* is reliable. It still does not follow that F raises the probability of G′. For assumptions (1) and (2) are compatible with the possibility that F will cause the subject to use *r rather than a more reliable process* (cf. Jäger forthcoming). Let *r*′ be the process the subject is likely to use given −F, that is, if G did not result from a reliable process. And let R′ state that the subject will use *r*′. To get the probability premise of CPS from (1) and (2), we also need (5):

(5) $P(G'/R\&F) > P(G'/R'\&{-}F)$.

Unless (5) is true, P(G′/G&F) may actually be lower than P(G′/G&−F).[12] Goldman & Olsson provide no reason for thinking that (5) is true, and it is not easy to find any.

The first assumption on which CPS is based, (1), is itself problematic. For it is quite possible that G results from two processes, one reliable and the other unreliable. Even if we ignore the possibility of overdetermination, G may have resulted from a reliable process that was the last stage of an overall unreliable process. Consider the following process for finding out the color of an apple.

u An unreliable apple-color belief-forming process.

(a) Toss a pair of dice and observe the result.
(b) If the dice fall snake eyes, look at the apple in good light, etc.;
    1. If the apple looks red, infer that it is red.
    2. If the apple looks green, infer that it is green;
    3. If the apple looks yellow, infer that it is yellow.

(c) If the dice fall any other way, look at a randomly selected photo of a pear in any light;
    1. If the photo of the pear looks red, infer that the apple is red;
    2. If the photo of the pear looks green, infer that the apple is green.
    3. If the photo of the pear looks yellow, infer that the apple is yellow.

[12] For an example, suppose that the safecracker in our pack of bank robbers is blind. In fact, he relied on Eddie to tell him when he is in the bank. Suppose further that if the safecracker had not relied on a reliable source, the leader of the pack would have instructed him to listen to Jimmy, who is even more reliable than Eddie. Jäger's (forthcoming) example involving different navigation systems has the same formal structure.

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

Process *u* is clearly less reliable than the processes people ordinarily use to determine the color of an apple. It can be expected to produce true beliefs at best a third of the time. Nevertheless, if the dice fall snake eyes, this unreliable process is likely to produce a true belief. For if it does, that belief will be produced by subprocess *u*(b), which is reliable (let the 'etc' be filled into make this true). If a subject uses *u* and *u*(b) in this way, and as a result forms the true belief that the apple is red, why should we assume that the subject will use subprocess *u*(b) again rather than the whole process *u*? If the subject reuses *u*(b), he is likely to get more true belief. If he reuses *u*, he is unlikely to get more true belief. The truth of G&F is compatible with both outcomes.

We might be able to determine which process the subject would reuse if we can assume that the subject is aware of the reliability of the processes producing his true beliefs (Jäger forthcoming). If the subject has suitable metabeliefs, and recognizes that *u*(b) is reliable while *u* is unreliable, then it is reasonable to assume that the subject will reuse *u*(b) but not *u*. While Olsson (2007, pp. 348, 352) and Goldman and Olsson (2009, p. 29) do countenance similar internalist conditions, Goldman has generally resisted them (see e.g., 1992, p. 434; 1999). So let us consider the externalist hypothesis that reuse of the relevant processes is "wired in" as part our innate equipment.

There are two ways we might interpret this externalist reuse hypothesis. On the first, what is wired in are certain processes such as sense-perception and memory. Such processes occur constantly and automatically from birth or shortly thereafter, much like breathing. *u*(b) is such a process, but not *u*. The fact that a true belief resulted from a process that is wired in is not necessary for it to be knowledge, however. Many inference rules mathematicians use, for example, are the product of much learning. Thus, resulting from a wired-in process is not necessary for knowledge to be more valuable than a mere true belief.

On the second interpretation of the externalist reuse hypothesis, what is wired in are not particular processes, but a disposition to use successful processes again. This would be something like a built-in inductive mechanism. There are several problems with this suggestion. The first is that success for the externalist is either the truth of the belief or its being produced by a reliable process (or both). The truth of a belief is not something the wired-in inductive mechanism would have an independent test for. Beliefs are produced internally by the belief-forming processes, but their truth is an external matter. A fortiori, being produced by a reliable process is also not something the wired-in inductive mechanism would have an independent test for. For a reliable process is defined as one that generally produces true beliefs. The hypothesized inductive mechanism is reminiscent of the conditioned learning processes studied by psychologists. But in those processes, the trigger for repeating an action is a perceived stimulus or reward of some sort, not the objective truth of a belief. Goldman and Olsson (2009, p. 29) seem to suggest that the trigger is finding the results "unobjectionable" and free of "apparent problems." These seem to be internal factors that are not very good criteria for the truth of the belief produced.

There is further trouble for the hypothesis that we somehow have a wired-in disposition to reuse processes that happen to produce a true belief. Many unreliable

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

processes will deliver a true belief on their first use, as *u* did above. What then keeps us from being condemned to continue using unreliable processes if they ever produce true beliefs? There must be a trigger for halting them. Presumably, that would be its production of a false belief. But reliable processes too will occasionally produce false beliefs. That cannot be enough to call a halt to a reliable procedure. So any wired-in disposition to reuse a process would have to be triggered not by a single success, but by a series of trials in which success is more frequent than failure. Such a disposition could not be invoked, however, to justify the conditional probability solution to the extra value of knowledge problem. The problem is to account for the value of a particular instance of knowledge. The solution assumes that given a true belief, the subject is likely to reuse the process that produced it. In terms of the wired-in hypothesis, a single success must trigger reuse (Jäger forthcoming).

The truth of both (5) and the probability premise of CPS depends on there being some probability that the subject will continue to be alive and well after forming the belief described by G. If the subject is certain to die, for example, both P(R/G&F) and P(R/G&−F) will equal 0. This implies that knowledge is no better than mere true belief in the subject's final moments.[13] While Goldman and Olsson (2009, p. 29) are willing to accept the consequence that the extra value of knowledge is contingent on empirical conditions, it is particularly hard to accept that although having a true belief is valuable for someone about to die, knowledge has no extra value. This would imply that it does not matter whether the individual's belief is justified or unjustified, which never seems true. If we compare two police officers making an arrest, one who knew that there was incriminating evidence before being shot dead and the other who merely had a true belief but survived, we would judge the former to have been in a better epistemic position and to have acted with greater propriety. Again, Goldman and Olsson seem at best to have shown that in the proper conditions, knowledge indicates a certain likelihood of additional value in the future. That falls short of accounting for the fact that knowledge is more valuable than mere true belief.

5 The linkage problem

Our question is why a true belief that is knowledge has more value than a mere true belief. The answer has to specify a property of that belief. The property specified by process reliabilism is *resulting from a reliable process*. The fact that the belief results from the process plays no role in the conditional probability solution, however. The prediction of future true beliefs is based entirely on the fact that a subject who has used a reliable process is likely to *use the process again*. Even if G&F has more value than G, it has no more value than a number of other conjunctions containing G.

13FL01 [13] As noted in footnote 7, Goldman and Olsson might observe that G&F would raise the probability of G′
13FL02 if the subject were going to survive. But this implies at most that the subject's knowledge *would be more*
13FL03 *valuable* than mere true belief in certain counterfactual conditions, not that it *is* more valuable.

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

(6) G&F has no more value than $G\&F_f$ or $G\&F_g$.

F_f: Some true belief resulted from a reliable belief-forming process.
F_g: Some reliable belief-forming process will be reused.

Given (6), we cannot infer from CPS that any property of G accounts for the extra value of its being knowledge.

Goldman and Olsson recognize this linkage problem.

> A compound state of affairs consisting in a reliable process followed by a true belief will be more valuable than the same true belief not preceded by a reliable process, and this is so even if there is no causal relation between the two, and hence no knowledge (Goldman and Olsson 2009, p. 34).

For a solution, they turn again to the motive analogy.

> We can simply note that our valuations are sensitive to causal linkages between suitable pairs of states. For example, good actions that are *caused* by good motives get higher moral marks than good actions that are merely *preceded* by good motives. The valuation of knowledge comports with this pattern. (Goldman and Olsson 2009, p. 34)

If Goldman and Olsson are going to use the motive analogy to defend reliabilism as a solution to the extra value of knowledge problem, they need to focus on the plausible fact that while a situation in which a good action is merely preceded by a good motive has some positive value, a situation in which the motive caused the action has more value. That would support, by analogy, the conclusion that having a true belief caused by a reliable process (knowledge) is more valuable than simply having a true belief and a reliable process. The pattern illustrated by the motive analogy is not very general, however. A situation in which Linda has a good espresso maker and a good cup of espresso does not obviously get lower marks than one in which the cup was produced by that espresso maker. And as Goldman and Olsson themselves note (thanks to Dennis Whitcomb), a situation in which one gets headache relief as a result of taking an aspirin is no better for the sufferer than a situation in which the headache relief is unrelated to taking the aspirin. A situation in which there is a causal relation between a man's intelligence and his winning a game of chance (both good things) may actually be made worse if the man won as a result of his intelligence (via cheating). Why should we assume that the process-belief case is like the unusual motive-action case?

One way the motive-action case differs is that it is good for a person to have a good motive even if the agent does not succeed in performing good actions. The value of a good motive is not completely instrumental. On a reliabilist theory, however, the value of a belief-producing process is entirely instrumental, as is that of aspirin. Goldman and Olsson address this problem by hypothesizing that a psychological process called "value autonomization" sometimes leads people to attribute intrinsic value to something that was originally assigned merely instrumental value. Even if we ignore the fact that their hypothesis is about value attribution rather than value, it remains ad hoc. Why should we assume that

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

autonomization has been applied to reliable belief-forming processes if it is postulated that the only fundamental epistemic value is truth?[14]

In the case of knowledge, it is not simply that a situation in which the subject truly believes something while also having the extra factor necessary for knowledge is better than a situation in which the subject truly believes it without that factor. The extra factor makes the belief better. That is, our believing the proposition is a better thing when the belief is knowledge and not just a true belief. This may be why Goldman and Olsson characterize the motive analogy by saying that good actions are better if they are caused by good motives. This characterization is dubious, however. To be sure, it is better for a surgeon to amputate a leg out of a desire to save the patient's life than for her to amputate the leg out of a desire to torture the patient or make money. But it does not follow that the goodness of the amputation depends on its cause. The pattern Goldman and Olsson thought the motive case illustrated fails quite generally with processes and products. Given two cups of espresso identical in taste, aroma, temperature, and all the other things that make for a good cup of espresso, the fact that one was produced by a good espresso maker will not make it a better cup of espresso than the other (Zagzebski 2000).[15] Most pertinently, the fact that a reliable belief-producing process turned out a true belief does not make the latter any better for the purpose of having true beliefs. We see again that the most we can conclude from a reliabilist account of knowledge is that knowledge is a better indicator of future true belief than mere true belief. The knowledge does not even have additional instrumental value. The reliable process has an instrumental value. The fact that the belief was produced by that process does not.

6 Concluding observations

Jones, Swinburne, Zagzebski and others have presented general considerations for thinking that a purely reliabilist epistemology cannot account for the extra value of knowledge. We have argued that Goldman and Olsson's attempt to overcome these considerations is unsuccessful. We also agree with Jones and Swinburne that the failure of process reliabilism in this endeavor shows that an analysis of knowledge as true belief produced by a reliable process is deficient. Specifically, it fails to imply that knowledge is justified belief. Goldman and Olsson claim that if certain plausible empirical assumptions are granted (such as that a reliable process is likely to be used again), the extra value of knowledge can be explained. The problem is that the empirical assumptions they cite do not ensure that the true beliefs in question are justified, and for that reason do not account for those beliefs having the extra value that knowledge has over mere true belief. The empirical conditions at best ensure that knowledge indicates additional true belief.

14FL01 [14] Cf. Goldman 2001, where he argues for "veritistic unitarianism," defined as the view that the core
14FL02 epistemic value is true belief.

15FL01 [15] See also Zagzebski 2004, 190 ff; Brady 2006; Brogaard 2006; Pritchard et al. 2010, Chap. 1.

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

We are dubious, however, that the extra value of knowledge problem has a "deep" solution. We set out the problem by observing the ways in which knowledge differs from true belief. One of the things knowledge requires, in addition to having a true belief, is that the subject's belief is justified. Justification is a positive normative property. For S to be justified in believing something is for it to be rational for S to believe it. If reasons are rationally required, S has good reasons for believing it: S reasoned properly and relied on good evidence. It is therefore tautological to observe that it is better for a true belief to be justified than unjustified. We can no more explain why it is better for our beliefs to be justified than we can explain why it is better for our actions to be moral, or our societies just.

This is not to say, of course, that it is obvious that knowledge always has more *practical* benefits than true beliefs. That may be as false as the parallel claim that acting morally always has more practical benefits than acting immorally. There are fundamentally different kinds of values, including the moral, the practical, and the epistemic or rational. These cannot be reduced to each other, or to nonnormative properties. If the extra value of knowledge problem is looking for such reductions, it should be dissolved.

References

Baehr, J. (2009a). Is there a value problem? In A. Haddock, A. Millar, & D. Pitchard (Eds.), *Epistemic value* (pp. 42–59). Oxford: Oxford University Press.

Baehr, J. (2009b). Inverting the value problem, paper presented at the International Conference "Epistemic Goodness", March 12–14, 2009, University of Oklahoma (unpublished).

BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge: Harvard University Press.

Brady, M. (2006). Appropriate attitudes and the value problem. *American Philosophical Quarterly, 43*, 91–99.

Brogaard, B. (2006). Can virtue reliabilism explain the value of knowledge? *Canadian Journal of Philosophy, 36*, 335–354.

Goldman, A. I. (1979). What is justified belief. In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: D. Reidel; reprinted in E. Sosa, J. Kim, J. Fantl, & M. McGrath (Eds.), *Epistemology: An anthology* (2nd ed., pp. 333–347), Oxford: Blackwell.

Goldman, A. I. (1986). *Epistemology and cognition*. Cambridge: Harvard University Press.

Goldman, A. I. (1992). Reliabilism. In J. Dancy & E. Sosa (Eds.), *A companion to epistemology* (pp. 433–436). Oxford: Blackwell.

Goldman, A. I. (1999). Internalism exposed. *The Journal of Philosophy 96*, 271–293; reprinted in Goldman, A. I. (2002). Pathways to knowledge (pp. 3–23). Oxford: Oxford University Press.

Goldman, A. I. (2001). The unity of the epistemic virtues. In A. Fairweather & L. Zagzebski (Eds.), *Virtue epistemology* (pp. 30–48), Oxford: Oxford University Press; reprinted in Goldman, A. I. (2002). Pathways to knowledge (pp. 51–72). Oxford: Oxford University Press.

Goldman, A. I., & Olsson, E. (2009). Reliabilism and the value of knowledge. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic value* (pp. 19–41). Oxford: Oxford University Press.

Greco, J. (2002). Virtues in epistemology. In P. K. Moser (Ed.), *The Oxford handbook of epistemology* (pp. 287–315). Oxford, New York: Oxford University Press.

Greco, J. (2003). Knowledge as credit for true belief. In L. Zagzebski & M. DePaul (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology* (pp. 111–134). Oxford: Oxford University Press.

Greco, J. (2010). *Achieving knowledge: A virtue theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.

Jäger, C. (forthcoming). Reliabilism and the value problem. *Theoria*

Jones, W. (1997). Why do we value knowledge? *American Philosophical Quarterly, 34*, 423–439.

Journal : Small-ext 11098 | Dispatch : 16-9-2010 | Pages : 13
Article No. : 9620 | LE | TYPESET
MS Code : | CP | DISK

Kvanvig, J. (1998). Why should inquiring minds want to know? *The Monist, 81*, 426–451.
Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese, 158*, 345–361.
Olsson, E. J. (2007). Reliabilism, stability, and the value of knowledge. *American Philosophical Quarterly, 44*, 343–355.
Plato: *Meno.*
Pritchard, D. (2007a). Recent work on epistemic value. *American Philosophical Quarterly, 44*, 85–110.
Pritchard, D. (2007b). The value of knowledge. In E. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (published August 2007) URL: http://plato.stanford.edu/entries/knowledge-value/.
Pritchard, D., Millar, A., & Haddock, A. (2010). *The nature and value of knowledge: Three investigations*. Oxford: Oxford University Press.
Riggs, W. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research, 64*, 79–96.
Sosa, E. (2003). The place of truth in epistemology. In L. Zagzebski & M. DePaul (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology* (pp. 155–180). Oxford: Oxford University Press.
Swinburne, R. (1999). *Providence and the problem of evil*. Oxford: Oxford University Press.
Swinburne, R. (2001). *Epistemic justification*. Oxford: Oxford University Press.
Weiner, M. (2009). Practical reasoning and the concept of knowledge. In H. Adrian, M. Alan, & P. Duncan (Eds.), *Epistemic value* (pp. 163–182). Oxford: Oxford University Press.
Zagzebski, L. (2000). From reliabilism to virtue epistemology. In G. Axtel (Ed.), *Knowledge, belief, and character* (pp. 113–122). New York: Rowman and Littlefield.
Zagzebski, L. (2003). The search for the source of the epistemic good. *Metaphilosophy, 34*, 12–28.
Zagzebski, L. (2004). Epistemic value monism. In J. Greco (Ed.), *Ernest Sosa and his critics* (pp. 190–198). Oxford: Blackwell.

Journal : **Small-ext 11098** Dispatch : **16-9-2010** Pages : **13**
Article No. : **9620** □ LE □ TYPESET
MS Code : ☑ CP ☑ DISK