

Reliabilism and the Value Problem

Christoph Jäger, Innsbruck

Draft May 2010

forthcoming in *Theoria* (2010)

Alvin Goldman and Erik Olsson (forthcoming) have recently proposed a novel solution to the value problem in epistemology, i.e., to the question of how to account for the apparent surplus value of knowledge over mere true belief. Their “conditional probability solution” maintains that even simple process reliabilism can account for the added value of knowledge, since forming true beliefs in a reliable way raises the objective probability that the subject will have more true belief of a similar kind in the future. I show that this proposal (i) confronts internal problems and (ii) implicitly invokes epistemic higher-level conditions that run against the spirit of externalism.

1. Introduction

Reliabilist theories of knowledge, it has been argued, cannot account for the fact that knowledge is generally more valuable than mere true belief. Reliabilism seems to suffer from the so-called “swamping problem”: the value of a true belief that has been generated in a reliable way (through reliable processes, faculties, or agents) appears to overwhelm the value of its reliabilist origin. For example, process reliabilism assigns positive value to reliable epistemic processes because such processes tend to produce true beliefs. However, how could the fact that a given item of knowledge has been produced in a reliable way add anything of axiological significance to its true-belief component? The cardinal aim in the realm of the epistemic is truth. When a subject knows that p , that aim has been attained; so the fact that this true belief was generated in a reliabilist way appears to be unable to confer any additional

value upon this belief. The value of the process seems to be entirely derivative from the value of having true belief, hence the fact that a given true belief constitutes reliabilist knowledge appears to be axiologically vacuous.¹

Before proceeding, let us distinguish this form of the value problem from other, partly related, topics debated in the literature. (i) First, the value question is not why the concept of knowledge is valuable. The concept might be valuable in social contexts in which it is beneficial to keep tabs on which information is correct and which informants are reliable. Ascribing knowledge may plausibly be regarded as an effective way to do so. This is an interesting idea, but it will not be discussed in this paper. (ii) Second, the value question is not simply why knowledge is valuable. A common answer to that question is that knowledge is at least true belief, and that true belief is valuable. An important issue here is whether this latter fact should be spelled out in terms of instrumental or, alternatively, intrinsic value. Here it will be assumed that true belief is at least instrumentally valuable: it enables us to reach our goals, and to do so in appropriate ways. Note that if true belief is only instrumentally valuable, beliefs about trivial subject matters (such as, in typical circumstances, the number of blades of grass on your lawn, etc.) are not valuable even if they are true. Similarly for knowledge. If there are worthless truths, it is hard to see how knowledge of such truths could be valuable at all.² If it is not, knowledge is a fortiori not generally more valuable than mere true belief. In what follows I shall thus construe the value question as concerning the greater value of at least many kinds of knowledge over mere true belief.

¹ For helpful expositions of this problem, see for example Jones (1997); Zagzebski (2000), (2004); Swinburne (1999, 2001); Kvanvig (2003); Sosa (2003); Brady (2006); Riggs (2002), (2007); Pritchard (2007a, 2007b); Greco (forthcoming), chapter 9. The label “swamping problem” is Kvanvig’s; Zagzebski (2004) and Brady (2006) call it the “value problem (in epistemology)”.

² For arguments to this effect see for example Sosa (2003), Baehr (forthcoming).

(iii) Third, some philosophers have formulated a more general comparative question. For example, in his booklength study of the value of knowledge Jonathan Kvanvig writes that explaining the value of knowledge “requires showing that knowledge is more valuable than any proper subset of its constituents” (p. 107). An even more demanding question Kvanvig poses is “whether and how knowledge has a value exceeding that of its parts” (p. x). This latter passage doesn’t talk about proper subsets, but the idea seems to be that the value of knowledge may even exceed the value of the sum of all its constituents. While this might be a theoretical option, I doubt that it corresponds to any widespread pre-theoretical intuition. In any event, in what follows I will not pursue the general questions Kvanvig raises. Instead, talk about the “value problem” or the “swamping problem” will refer to the moderate comparative question of why at least many kinds of knowledge are at least instrumentally more valuable than the true beliefs they involve.

Alvin Goldman and Eric Olsson (forthcoming) have proposed an interesting and novel solution to the value problem. In fact, they argue, even simple process reliabilism can account for the extra value of knowledge. Simple process reliabilism is the view that S knows that p if and only if (i) p; (ii) S believes that p; and (iii) S’s belief that p was generated through a reliable process.³ Typically, process reliabilists add some anti-Gettier clause. As Goldman and Olsson note, however, “the idea of knowledge depending on the existence of a reliable connection is the central one behind reliabilism, and it would be unfortunate for the theory if that very component failed to produce an added value” (p. 6). Indeed, so let us focus on the question whether simple process reliabilism can solve the swamping problem along the lines Goldman and Olsson propose.

The additional valuable property which, according to Goldman and Olsson, distinguishes (simple process) reliabilist knowledge from mere true belief is the property of

³ Cf. Goldman & Olsson forthcoming, p. 4. (Page references here and henceforth to the manuscript.)

raising the conditional objective probability that the subject's "future beliefs of a similar kind" will be true as well. Introducing a diachronic perspective on the truth goal of believing may point in a promising direction. However, Goldman and Olsson have a skeleton in their closet. As it stands, their conditional probability solution (henceforth: CPS) is committed to the claim that S's knowing that p exceeds the value of S's merely truly believing that p only if S believes that S truly and reliably believes that p. Introducing such higher-level epistemic conditions however sits ill with the general spirit of externalism.

2. The conditional probability solution

CPS proposes that on (simple process) reliabilism,

“the [objective] probability of having more true belief (of a similar kind) in the future is greater conditional on S's knowing that p than on S's merely truly believing that p” (p. 11).

Why would this be so? In support of their thesis, Goldman and Olsson present a (slightly modernized) version of Plato's famous Larissa-example from the Meno. Suppose you are driving to Larissa, and there are two forks on the way. Situation 1: You are equipped with a reliable onboard navigation system which, when you reach the first crossroads, tells you correctly that the shortest route to Larissa is to the right. Situation 2: Your onboard navigation system is unreliable; yet when you reach the first crossroads it also recommends correctly that you take a right. “On the simple reliabilist account of knowledge”, Goldman and Olsson argue,

“you have knowledge that p in Situation 1 but not in Situation 2. This difference also makes Situation 1 a more valuable situation (state of affairs) than Situation 2. The

reason is that the conditional probability of getting the correct information at the second crossroads is greater conditional on the navigation system being reliable than conditional on the navigation system being unreliable” (pp. 12f.).

So the idea is that having a reliably produced true belief B raises the conditional probability that future beliefs of a similar kind will be true as well, because forming B raises the relative probability (i.e., relative to using a different process from the one that produced B) that such future beliefs will be produced by the same (type of) reliable process that generated B.⁴ The crucial question is: Why would this be so? This question divides into two parts. (i) Why would an epistemic subject re-employ a reliable epistemic mechanism? (ii) Would employment of such a mechanism in the future indeed raise the conditional probability of having more true belief on those future occasions, compared to a situation in which the mechanism does not operate in the original circumstance? Let us begin with this second question.

An initial worry is that nothing in Goldman and Olsson’s account rules out that, if some reliable epistemic mechanism R were not operating in a given belief forming process, this might cause a still more reliable mechanism R* to operate, not on the present, but on the relevant future occasions.⁵ In that case, it is more probable that S acquires more true belief on those future occasions if on the present occasion S does not employ R but acquires – for example by sheer guesswork – a true belief that has not been produced by any reliable mechanism. Suppose that on your way to Larissa you are using a navigation system N1

⁴ It is not clear whether this is supposed to hold for all, most, many, or maybe just some future beliefs of a similar kind. A moderate version of the extra value claim would already be supported by the fact that for some future beliefs the probability that they will be true is raised. The argument presented here applies not only to this construal of CPS but also to the stronger ones.

⁵ I owe this point to Wayne Davis.

which, by suitable reliabilist standards, counts as reliable. Nevertheless, let us suppose that N1 gives you the correct route at the first and second crossroads, but not at the following three. Suppose further that, had you not been using this system at the first crossroads but instead just guessed correctly where to go (for example, because you believed the system was broken), you would have stopped at the next service station and bought a top quality system N2 that is much more reliable than N1. In that case, the probability that you acquire more true geographical beliefs for the rest of your trip is greater than in a situation in which you use the low quality system N1 at the first crossroads, and continue to use it. So employing the reliable mechanism N1 does not raise but instead lowers the conditional probability that you acquire more true beliefs of a similar kind in the future.

The second question (i.e., question (i) mentioned above) is: Exactly why, and under what conditions, would one re-employ a reliable epistemic process or mechanism? Goldman and Olsson argue that “the extent to which a knowledge state enhances the conditional probability of future true beliefs depends on a number of empirical regularities” (p. 12). They mention four such regularities. (i) The first is “Non-uniqueness”, which refers to the fact that the epistemic subject is likely to encounter problems of the same type on future occasions as well. (ii) Next there is “Cross-temporal access”, which means that processes or methods that have proven to be reliable will also be available at future times. (iii) “Generality” is shorthand for the fact that it is likely for a given method that was reliable on one occasion also to be reliable in other, similar situations. Finally, (iv) a “learning” condition (for short: Learning) must be fulfilled, to the effect that,

“if you have used a given method before and the result has been unobjectionable, you are likely to use it again on a similar occasion, if it is available. Having invoked the navigation system once without apparent problems, you have reason to believe that it

should work again. Hence, you decide to rely on it also at the second crossroads” (p. 13).

So, Goldman and Olsson’s response is that \underline{S} will re-employ a reliable epistemic process or mechanism if certain results it produced on other occasions are “unobjectionable” for \underline{S} and \underline{S} has therefore “reason to believe that it [the process or mechanism] should work again”. (They call the learning condition “empirical”, I take it, because it is a contingent fact about actual epistemic subjects that they – sometimes – learn from their epistemic performances.)

Before looking more deeply into the implications of these empirical regularities, let us try to provide a somewhat more precise description of the problem to which CPS responds. As already indicated, there is more than one question that philosophers pursue under the label “value problem” (in epistemology). Note that, whatever the boons and banes of CPS, if it invokes “empirical regularities”, the solution in any case will not satisfy those critics of reliabilism who believe that, if knowledge is more valuable than mere true belief, the analysis of knowledge alone must be in a position to explain this. Goldman and Olsson (pp. 7-11) extensively discuss Jonathan Kvanvig’s set up of the question (in Kvanvig 2003). Besides the fact that Kvanig sets himself the more demanding task of showing that the value of knowledge exceeds the value of any proper subset of its parts, he contends that

“a satisfactory answer to the question of the value of knowledge will need to explain why knowledge is, by its very nature, more valuable than its parts. It will not be enough, for example, to show that sometimes or in some places knowledge is more valuable than its parts. Instead, we will need to show that no matter what the world happens to be like, knowledge is more valuable” (Kvanvig 2003, xiv, emphasis C.J.).

Similarly, according to Wayne Riggs philosophers have traditionally been committed to the principle that knowledge “is always more valuable than (mere) true belief”. Hence, Riggs argues, “a condition of adequacy on a complete theory of knowledge is that it entail (and preferably explain)” this value principle (Riggs 2002, 79). In the same vein Jason Baehr has more recently portrayed what he calls the “standard conception” of the value problem as a conception that “is thought to motivate a constraint on an analysis of knowledge such that any plausible analysis must entail that knowledge is more valuable than true belief” (Baehr forthcoming, 1).⁶ So according to this conception the added value claim holds that:

(AV_S) Necessarily, knowledge is more valuable than mere true belief.

Call this the “strong version” of the added value claim. Although Goldman and Olsson sometimes present CPS as if it included an answer to Kvanvig and Riggs, clearly CPS does not provide an answer to (AV_S). However, another, weaker version of the added value claim merely says that:

(AV_M) If certain contingent conditions are satisfied, the value of knowledge exceeds the value of mere true belief.⁷

Call this the “moderate version” of the added value claim. CPS attempts to show that, contrary to what many epistemologists have argued in recent years, process reliabilism has the

⁶ Baehr however criticizes and eventually rejects this conception.

⁷ This is the position adopted for example by Jones (1997).

resources to explain why (AV_M) is true, i.e., why in the actual world knowledge is more valuable than mere true belief.⁸

Let us now take a closer look at the “empirical conditions” under which CPS is said to work. As quoted above, the idea is that under certain circumstances reliable epistemic mechanisms tend to be re-employed because the subject believes that they “should work again” and hence “decides” to rely on them on future occasions as well. What does such a decision require? Concerning Non-uniqueness, it would seem to require not only the fact that a problem of the type in question has reoccurred, but also that the subject is aware of this fact. Similarly for Cross-temporal Access: Not only must the method in question be available to you. A reason-based decision to use it again will also rest on the belief that it is still available. Regarding Generality, note that not only must it be likely that the method was reliable on past occasions and that it will be reliable again, but that a decision to re-employ it will also be based upon your belief that the method is reliable in situations of the type at hand.

⁸ “On this approach”, Goldman and Olsson write about CPS, “the greater value is attained normally, not always” (p.1). However, the authors believe that people nevertheless “tend to think that knowledge is always more valuable” (ibid., emphasis C.J.). Therefore, a second step in Goldman and Olsson’s project is to explain this latter phenomenon. They suggest “a general psychological process of value autonomization whereby a reliable process comes to be regarded as valuable in itself as opposed to being valuable only in virtue of the true beliefs it causes” (p.1). In conversation, the authors told me that this latter topic of their paper was primarily developed by Goldman, while CPS was rather Olsson’s contribution. In my view this changes little for Goldman’s commitments to CPS and the higher-level epistemic conditions this account involves. For the theory of value autonomization he presents in the paper is explicitly motivated by the question why people widely think that knowledge is always more valuable than mere true belief, when it is in fact only sometimes more valuable – as it is explained by CPS. I think it is doubtful that people widely regard knowledge of whatever kind always to be more valuable than mere true belief. (Upon little reflection, for instance, many people realize that there is trivial knowledge, such as knowledge of how many

Finally, consider Learning, the condition according to which “you are likely to use ... [a given mechanism or method] again on a similar occasion”. What does it mean for an epistemic mechanism or method that has been employed in a given situation that “the result is unobjectionable” for the subject and that the subject has “reason to believe that it [the mechanism] should work again”? First, the results in question are beliefs, and a belief is unobjectionable (from an epistemic point of view) for S only if S believes that the belief is true. Second, even though merely having a reason to believe something need not be a reason for which one actually does believe it, it is clear that, if one “has reason to believe” that a given mechanism or method should work again and on that basis decides to re-employ it, that reason also becomes effective, in the sense that one actually does believe on its basis that the mechanism will work again.⁹ But if you believe that some process or method should work again and hence decide to reuse it, you believe that it was, and still is, sufficiently reliable. We thus get the following result: CPS, Goldman and Olsson concede, works only if the Learning condition is satisfied, but Learning presupposes that S believes (i) that the mechanism employed in a given belief forming situation delivered a true belief, and (ii) that that process or method was, and still is, reliable. Hence on Goldman and Olsson’s account acquiring reliabilist knowledge on a given occasion has more value for S than merely acquiring a corresponding true belief, only if S believes that, on that occasion, S has acquired a true belief

blades of grass there are on your lawn.) However, here I don’t have the space to discuss Goldman’s account of value autonomization, but shall instead confine myself to CPS.

⁹ The Learning Condition also figures in Olsson (2007), where he describes it as follows: “A method that was unproblematically employed once will tend to be employed again on similar problems in the future” (p. 348). And in a footnote Olsson explains that, as he uses the notion, “a belief generating method was ‘unproblematically employed’ if there is no positive sign suggesting to the believer that the belief thus produced might be false” (p. 353, emphasis C.J.).

that was produced by a reliable mechanism. It thus emerges that CPS relies on the following doxastic higher-order principle:

$$(P) \quad \Box [(RB_{Sp} \ \& \ p \ \& \ (V(RB_{Sp} \ \& \ p) > V(B_{Sp} \ \& \ p \ \& \ \sim RB_{Sp})) \supset B_S(RB_{Sp} \ \& \ p)]$$

Necessarily: If S reliably and truly believes that p, and the value of this fact exceeds the value of S's merely truly believing that p, then S believes that S reliably and truly believes that p.¹⁰

Some comments are in order. First, although Learning is an “empirical regularity” that does not hold in every possible world, CPS maintains that Learning is a necessary condition of the added value of reliabilist knowledge. So the argument just laid out may also be put by saying that according to CPS there is no possible world in which reliabilist knowledge has that added value but S does not believe of him/herself that he or she reliably and truly believes that p. Second, prima facie it may appear as if S's believing the second conjunct of the belief content stated in the consequent (i.e., p) is just a trivial a fortiori consequence of the fact that, as stated in the antecedent, S reliably believes that p. However, it is highly controversial whether belief is closed under conjunction introduction; in fact there are strong reasons for thinking that it is not. So even though S's reliably believing that p does of course entail that S believes that p, it does not follow that S, if S also believes that he or she reliably believes that p, believes the conjunction ($RB_{Sp} \ \& \ p$).

Finally, on simple process reliabilism S's having a reliably formed true belief that p just means that S knows that p. It may thus be argued that (P) amounts to saying that:

¹⁰ Strictly speaking, this way of stating the consequent is not entirely accurate. What is required is that S believe him/herself to know that p. This may not be the case if S refers to him- or herself via a proper name or a definite description. We can ignore this complication for present purposes.

$$(P^*) \quad \Box [(K_{Sp} \& (V(K_{Sp}) > V(B_{Sp} \& p \& \sim K_{Sp}))) \supset B_S(K_{Sp})]$$

Necessarily: If \underline{S} (reliably) knows that p , and the value of this fact exceeds the value of \underline{S} 's merely truly believing that p , then \underline{S} believes that \underline{S} (reliably) knows that p .

The problem with (P*) is that the subject need not believe that a reliably formed true belief constitutes knowledge. Not everyone is a reliabilist. Hence if the consequent of (P*) is construed as a de dicto belief ascription according to which \underline{S} believes what he or she could state, in English, by saying “I know that p ”, (P*) is not generally true.¹¹ However, what this reflection shows is that CPS does imply (P*) for subjects who are reliabilists, i.e. for those of us who do believe that a reliably formed true belief constitutes knowledge. In their paper, Goldman and Olsson also discuss, and defend, a notion of “weak knowledge” according to which in some contexts “knowing that p ” just means “truly believing that p ” (pp. 1-3). In such contexts knowledge trivially fails to be more valuable than mere true belief. Here I shall not discuss the argument for weak knowledge.¹² Yet it may be noted that, cast in terms of the notion of knowing weakly (knowing_w), my argument can also be put by saying that for reliabilists, and in contexts in which a given item of reliably generated first-order knowledge (knowledge_R) is supposed to be more valuable than mere true belief, CPS adopts what may be called a weak version of the KK-principle. In Goldman and Olsson’s account, if \underline{S} is a reliabilist, \underline{S} 's knowing_R that p is more valuable than \underline{S} 's merely truly believing that p , only if \underline{S} at least knows_w that S knows_R that p .

In another paper Olsson argues, in the context of a related argument, that the extra value of reliabilist knowledge derives, among other things, from the fact that inquirers are

¹¹ Thanks to Wayne Davis for alerting me to this point.

¹² I do critically discuss their argument in Jäger (forthcoming).

“track keepers in the sense that they keep a record of the sources of their belief” (Olsson 2007, p. 348). However, track-keeping, Olsson explicitly concedes, “is a modest internalist requirement on a cognitive agent. It requires that the agent maintain a mental record ... of how beliefs were acquired” (p. 352). Olsson then maintains with respect to this condition that, even though knowledge may best be defined in an externalist fashion, “the full realization of its value requires the satisfaction of a modest internalist condition” (ibid.). My argument shows that a similar kind of track-keeping is needed to explain why, as CPS maintains, epistemic agents reuse reliable epistemic mechanisms. Olsson’s overall argument in his 2007 paper for the extra value of reliabilist knowledge is that this kind of knowledge is conducive to successfully completing actions over time. The reason, he argues, is that reliably formed true beliefs are more stable than those which have been formed in an unreliable fashion. Suppose this is right. Although Goldman and Olsson don’t invoke this stability thesis in their paper, it would nicely complement CPS. For having “more true belief in the future” would seem to be not only a consequence of reusing epistemic processes or mechanisms and thereby produce new true beliefs; instead, it is also a function of giving the true beliefs one has acquired a stable position in one’s doxastic system. According to this view, there are thus two ways in which internalist constraints figure in an explanation of the extra value of knowledge. The Learning condition ensures that people re-employ reliable epistemic mechanisms (and thereby raise the probability of having more true belief of a similar kind in the future). The Track-keeping condition secures the diachronic stability of true beliefs that have been formed in a reliable way.

3. Wired in is no way out

A crucial question at this point is whether Goldman and Olsson could not abandon their conditions and replace them by certain externalist-friendly “empirical regularities”. One move that springs to mind is to replace the current version of the Learning Condition with the claim

that people are simply hardwired to reuse epistemic processes that have proven to be epistemically “unobjectionable”.¹³ Such a wired-in disposition would still constitute some kind of second-order epistemic mechanism. Yet invoking it could avoid (for example) access internalist constraints by not requiring that S have introspective doxastic access to the fact that she is equipped with such a mechanism. However, a mechanism M of the desired kind would have to be able to distinguish reliable epistemic processes from unreliable ones. Why else should it foster re-employment of the former rather than the latter? A crucial question is thus how M could accomplish this (in a non-internalist way). It will obviously not do to postulate that M is somehow able to recognize when a given belief is true. (For the mere fact that a belief is true, is of course compatible with its having been produced by an unreliable method or process.) In order for M to identify a method as reliable, M would apparently have to rely, not on a single success, but on a favourable track record of employments of that process or method. M would have to contain some kind of unconscious inductive mechanism that would generate a tendency or disposition in the subject to reuse epistemic processes which over a representative series of trials it records as successful, while at the same time blocking tendencies to reuse processes with an unfavourable track record.

However, the prospects for developing a convincing proposal along such lines do not appear rosy. First, I do not know of any evidence that such a mechanism actually exists. Second, even if there were such evidence, the value problem would still haunt us. For what is at issue is the question of how to account for the extra value of an arbitrary, single instance of (simple process) reliabilist knowledge. The value problem, as outlined above, is not how to account for the fact that a series of true beliefs produced by the same kind of reliable mechanism and of which the subject knows that they are true and reliable is more valuable than true beliefs that are not reliably produced. The question is how to account for the fact that any particular belief B, if it is produced in a reliable way, is more valuable than any non-

¹³ Goldman suggested this move in conversation.

reliable belief with the same content. Appeal to wired-in mechanisms that analyze and statistically evaluate doxastic track records falls short of telling us (i) why a single success of an arbitrary reliable epistemic process or method might have special value, and (ii) why this single performance of the mechanism would trigger re-employments on suitable future occasions. On the wired-in proposal, what is more valuable than generating a mere true belief is not generating a corresponding reliably formed true belief. Instead what has added value is generating several, presumably many reliably formed true beliefs of a similar kind. In this case the added value accrues to a whole series of reliably formed true beliefs, so the object of the added-value ascription has changed.

4. Objections, clarifications, conclusion

Although this should be clear from the preceding sections, let it be emphasized once more that my argument neither maintains nor entails that the concept of knowledge discussed in CPS degenerates into an internalist notion. Someone may raise the following objection. “As long as the internalist component is not a constituent of knowledge, the reliabilism and externalism of Goldman and Olsson’s position and response to the value problem remain intact”.¹⁴ This is half true, but half false. What is true is that nothing in my argument attacks – or was ever intended to attack – Goldman and Olsson’s reliabilist and externalist position with respect to the concept of knowledge. Still, what my argument shows is that by invoking epistemic higher-order constraints their response to the value problem leaves externalist grounds. CPS maintains that in all possible worlds in which the value of a reliably formed true belief that *p* exceeds the value of a corresponding mere true belief that *p*, the subject believes that he or she reliably and truly believes that *p*.

In this context, it is also worth emphasizing that, as an anonymous referee put it, “although Goldman may be reluctant to embrace that consequence of their response ..., Olsson

explicitly acknowledges ‘a modest internalist requirement on a cognitive agent’ (Olsson 2007, 352 ...) that is not part of the analysis of knowledge in terms of reliability but instead figures as an ‘essential element of the environment in which knowledge attains its maximum worth’ (ibid.)”. That is true (and I do quote Olsson on this in the concluding passages of section 2). But what is the objection? “The author [of the present paper, C.J.] notes this concession”, the critic goes on, “but does not see that it threatens to undermine his/her objection.” Well, the dialectic, and the relevance of my argument, are slightly more complicated. One should indeed distinguish its impact for Goldman and for Olsson, respectively. Regarding the former, it certainly is a relevant observation that the solution to the value problem Goldman officially adopts in the paper¹⁵ relies on higher-order epistemic conditions. A core motivation of reliabilism – and who has taught us this better than Goldman himself? – is to provide a non-Cartesian epistemology. But if one’s epistemology should be able to solve the value problem, my argument shows that the way in which process reliabilism claims to do so does not remain strictly non-Cartesian. For it must balance its general externalist orientation with higher-order epistemic constraints.

¹⁴ This objection was raised, in these words, by an anonymous referee.

¹⁵ Even if CPS is, as the authors told me in conversation, primarily an offspring of Olsson’s, it would be odd to assume that Goldman does not endorse a core part of a paper that bears his name as a co-author. Moreover, the account of value autonomization which (as the authors told me) has primarily been developed by Goldman, is explicitly motivated by the fact that, while “the conditional probability solution explains why reliabilist knowledge is normally but not always more valuable than mere true belief, ... some philosophers think we are always prepared to attribute greater value to knowledge than to mere true belief” (p. 15). This would seem to imply that Goldman does endorse CPS as well. Goldman’s theory of value autonomization, in other words, is not an independent theory. It starts from the assumption that CPS is true and that precisely this fact calls for an explanation of why people tend to think that knowledge is not only sometimes, but always more valuable than mere true belief. Apart from this, were Goldman not to endorse CPS, he would still owe us a solution to the value problem.

My answer with respect to Olsson is twofold. First, in light of what he says in Olsson (2007), the internalism of CPS indeed may not be regarded as a problem for his overall position. However, my argument brings out the precise way in which higher-order constraints enter into those “empirical regularities” he adopts in his and Goldman’s piece. I do regard this as a constructive contribution to Olsson’s position. Yet the problem of CPS discussed in the first part of section 2 remains for Olsson as well: There are counterexamples to the claim that using a reliable mechanism, even under the empirical conditions CPS considers, really does raise the probability of having more true beliefs of a similar kind in the future. Witness the example of the two navigation systems N1 and N2, where using N1 on a given occasion excludes using the much better system N2 on suitable future occasions, whereas not using N1 would result in using N2 in the future.

One final comment on internalism vs. externalism and their relevance for the analysis of knowledge in its relation to the value problem. Note that for the strong version of the added value claim (AV_S), (certain forms of) internalism would appear to make the grade if one were to adopt a solution in the spirit of CPS. For if S’s knowledge conceptually entails that S believes that S knows (i.e., believes that S has an epistemically adequate true belief), then there is no possible world in which S knows but fails to have the beliefs necessary for deciding to re-employ the epistemic mechanism or method S deemed successful. There is no space here, however, further to pursue the question whether (AV_S) is ultimately defensible or whether it should be replaced by the moderate version of the added value claim (AV_M).

With respect to reliabilist concepts of knowledge Goldman once asked: “Doesn’t the agent have to ... know that those processes [which cause true beliefs in the epistemic agent] are reliable? Or doesn’t he have to be ‘in a position’ to know these things (or at least believe them reliably)? The reliabilist replies that an added requirement of this sort is inappropriate” (1992, p. 434). The arguments presented in this paper show that with respect to CPS, as an answer to the axiological question about the surplus value of knowledge, an added

requirement of precisely this sort is not only appropriate but indeed follows from Goldman and Olsson's proposal. In "Internalism Exposed" Goldman writes that he sees "no hope for internalism. It does not survive the glare of the spotlight" (2002, p. 293). If what I have argued is on target, there is little hope for (simple process) reliabilism to handle the value problem in the way Goldman and Olsson propose unless knowledge is placed under the spotlight of the subject's mind.¹⁶

Literature

- Baehr, Jason (forthcoming): Is There a Value Problem?, forthcoming in Epistemic Value, ed. Duncan Pritchard, Alan Millar, & Adrian Haddock, Oxford: Oxford University Press, 2009
- Brady, Michael, 2006: Appropriate Attitudes and the Value Problem, American Philosophical Quarterly 43, 91-99
- Goldman, Alvin & Erik Olsson, forthcoming: Reliabilism and the Value of Knowledge, forthcoming in Epistemic Value, ed. Duncan Pritchard, Alan Millar, & Adrian Haddock, Oxford: Oxford University Press, 2009
- Goldman, Alvin, 2002: Internalism Exposed, reprinted in id., Pathways to Knowledge, Oxford: Oxford University Press, 3-23
- Goldman, Alvin, 1992: Reliabilism, in A Companion to Epistemology, ed. Jonathan Dancy & Ernest Sosa, Oxford: Blackwell, 433-36

¹⁶ Central parts of this paper were presented at the conference "Reliable Knowledge & Social Epistemology – The Philosophy of Alvin Goldman", organized by Gerhard Schurz and Markus Werning, May 2008 in Düsseldorf. For valuable comments, either in discussions on that conference or in conversations at other opportunities, I am indebted to Elke Brendel, Wayne Davis, Evan Fales, Alvin Goldman, Michael Gorman, Thomas Grundmann, Katherine Munn, Erik Olsson, Christian Piller, Gerhard Schurz, Markus Werning, and two anonymous referees for "Theoria".

- Greco, John, forthcoming: Achieving Knowledge, Cambridge University Press, 2010
- Jäger, Christoph, forthcoming: Why to Believe Weakly in Weak Knowledge, Grazer Philosophische Studien 79
- Jones, Ward, 1997: Why Do We Value Knowledge?, American Philosophical Quarterly 34, 423-39
- Kvanvig, Jonathan L, 2003: The Value of Knowledge and the Pursuit of Understanding, Cambridge, New York: Cambridge University Press
- Olsson, Erik J., 2007: Reliabilism, Stability, and the Value of Knowledge, American Philosophical Quarterly 44, 343-355
- Pritchard, D., 2007a: Recent Work on Epistemic Value, American Philosophical Quarterly 44, 85-110
- Pritchard, D., 2007b: The Value of Knowledge, in: The Stanford Encyclopedia of Philosophy, ed. Edward Zalta, URL: <http://plato.stanford.edu/entries/knowledge-value/> (published August 2007)
- Riggs, W., 2002: Reliability and the Value of Knowledge, Philosophy and Phenomenological Research (64), 79-96
- Riggs, W., 2007: The Value Turn in Epistemology, in New Waves in Epistemology, ed. Vincent Hendricks & Duncan Pritchard, Palgrave
- Sosa, Ernest, 2003: The Place of Truth in Epistemology, in Intellectual Virtue: Perspectives from Ethics and Epistemology, ed. Linda Zagzebski & Michael DePaul, Oxford: Oxford University Press, pp. 155-179
- Swinburne, Richard, 1999: Providence and the Problem of Evil, Oxford: Oxford University Press
- Swinburne, Richard, 2001: Epistemic Justification, Oxford: Oxford University Press
- Zagzebski, Linda, 2000: From Reliabilism to Virtue Epistemology, in Knowledge, Belief, and Character, ed. Guy Axtell, Rowman and Littlefield

Zagzebski, Linda, 2004: Epistemic Value Monism, in Ernest Sosa and His Critics, ed. John Greco, Oxford: Blackwell, 190-198