

Reliabilism, Veritism, and Epistemic Consequentialism

Alvin I. Goldman

To appear in the Tenth Anniversary Issue of EPISTEME (forthcoming)

ABSTRACT

According to Selim Berker the prevalence of consequentialism in contemporary epistemology rivals its prevalence in contemporary ethics. Similarly, and more to the point, Berker finds epistemic consequentialism, epitomized by process reliabilism, to be as misguided and problematic as ethical consequentialism. This paper shows how Berker misconstrues process reliabilism and fails to pinpoint any new or substantial defects in it.

1. THE ETHICAL AND EPISTEMOLOGICAL BACKGROUND

At some point in the distant past, ethics and epistemology were positioned in opposite corners of the philosophical establishment. One was the epitome of "theoretical" philosophy and the other the epitome of "practical" philosophy. Relatively little contact was made between them. Today, however, analogies abound between ethics and epistemology. It is widely accepted, for example, that epistemology has a strong strand of normativity and should therefore be viewed (at least in certain respects) as a legitimate cousin of ethics. Ethics has been the chief battleground of powerful debates between consequentialism and deontology, but epistemology has also witnessed an analogous set of disputes.

Selim Berker is immersed in both of these disputes. In two recent articles (Berker 2013a, 2013b) he highlights the pervasiveness of consequentialism, or teleology, in contemporary epistemology. The main form of epistemic consequentialism regards truth -- or true belief plus error avoidance -- as the "final" value or end in the epistemological domain. Epistemic teleologists take other epistemic goods to have derivative value in virtue of their instrumental relationship to the aforementioned final end or ends. Berker himself rejects all forms of epistemic consequentialism, particularly insofar as they

purport to explain justification, reasonability, rationality, etc. in terms of those truth-linked values or ends. He regards process reliabilism as "the paradigm example of a teleological epistemic theory" (2013a, 350), and it is the central target of his critique. This paper responds to Berker's critique. As the originator and long-time proponent of process reliabilism (see especially Goldman, 1979, 1986, 1992, 2012), I ask the reader's indulgence in my focusing attention on this particular theory of justification in the present response to Berker's considerably broader critique.¹ During the course of the paper I shall also address the connection between process reliabilism and a related, but distinct, epistemic theory, i.e., "veritism," which is also a target of Berker's critique.

Let us set the stage for Berker's arguments by considering some of the scholarly background he provides for the debate between consequentialism and deontology within ethics of the last century. Berker clearly sees the contemporary issues for epistemology through the lens of the ethics literature. This is important partly because the parallel with ethics leads him, in my opinion, to some misconstruals of reliabilism, which might not have been made if he were not so consumed with ethics parallels. Here are some of the particularly salient passages from ethicists that Berker cites in the run-up to his discussion of epistemic consequentialism. These quotations can help us understand where Berker is "coming from" in his interpretation of epistemic consequentialism.

The continuing debate between Kantianism and utilitarianism (the most prominent form of moral consequentialism) is vividly depicted by Christine Korsgaard in the following passage:

To later generations much of the moral philosophy of the 20th century will look like a struggle to escape utilitarianism. We seem to succeed in disproving one utilitarian doctrine, only to find ourselves caught in the grip of another.
(Korsgaard 1993, 275).

A similar passage by Philippa Foot attests to the depth and power of the dispute:

¹ It is broader in the sense that it takes aim at a wide swath of epistemologists, including internalists as well as externalists. There isn't enough space here to canvass Berker's treatment of all the others whom he regards as (equally misguided) epistemic teleologists or consequentialists.

It is remarkable how utilitarianism tends to haunt even those of us who will not believe in it. It is as if we forever feel that it must be right, even though we insist that it is wrong. (Foot 1985, 196).

Turning from moral to epistemic consequentialism, Berker himself writes:

I believe that a similar specter haunts contemporary epistemology: the specter of reliabilism. This essay represents my attempt at an exorcism. (Berker, 2013b, 363)

I see Berker's invocation of Korsgaard's and Foot's longstanding battle with utilitarianism as providing assistance in conjuring up a suitably fierce and persistent opponent for his battle with epistemic consequentialism. If, switching metaphors a bit, he can portray reliabilism as a suitably fierce dragon, slaying that opponent will be all the more worthy and creditable. I don't mind his portraying reliabilism as a battle-worthy dragon. If it is going to be portrayed as a dragon, however, I would like him to be straight on the sort of fire it breathes. A large part of the argument of the present paper is that Berker brings the wrong kind of fire-retardant to the task because he misconstrues the dragon's type of fire. Or, returning to Berker's own metaphor, he proposes misguided formulas to utter in trying to exorcise the specter because he misunderstands the nature of the spook.

Reliabilism, says Berker, might be better called "truth-conducivism." The guiding thought is that what makes a (token) belief epistemically justified is its connection to a process (or indicator, or method) that conduces toward the promotion of true belief and the avoidance of false belief. This is analogous to various kinds of ethical consequentialism or teleology, in which an action, rule, motive, etc. qualifies as morally right if and only if it conduces to good consequences, or better consequences than its alternatives. In the first and longer of the two papers Berker suggests that a principal shortcoming of this kind of theory is that it ignores the "separateness of propositions." This is a good place to begin, especially in light of the fact that the phrase "the separateness of propositions" appears in the title of this paper. What does this mean?

Consider Judith Thomson's (1976) example (similar to ones found in Foot 1978) of a transplant surgeon who encounters a patient during a routine check-up and decides to chop him up, take his organs, and transplant them in five other patients, each in need of one of those organs. From a teleological point of view, this action is good, perhaps even right. Greater overall happiness will be promoted by this action than by simply attending to the original patient and letting those in need of organ transplants die. Following many moral theorists, however, Berker argues that trade-offs between persons should in general be avoided. This point was initially introduced by John Rawls (1971), who rejected any principle of utility maximization for this kind of reason.

2. ALLEGED EPISTEMIC ANALOGUES OF ETHICAL TRADE-OFF PROBLEMS

Now consider epistemological analogues of such a trade-off cases. Berker presents a series of such examples, starting with one in the style of Roderick Firth (1981), though in a form shaped by Richard Fumerton (2001). In this example Berker (2013, 364) imagines himself to be scientist seeking a grant from a religious organization that only gives grants to believers. He is a longstanding atheist, but knows he couldn't lie about this matter. Driven by a keen desire to obtain many true scientific beliefs by means of winning the grant, he manages to form a belief in God's existence. Is this belief justified? Intuitively, the correct answer is "No". What verdict does epistemic, or veritistic, consequentialism deliver? Here is what Berker says:

From veritism's perspective, my coming to believe in God's existence would be a sacrifice for the greater (epistemic) good: I would be trading off the pursuit of one thing of final epistemic value (true belief and the avoidance of false belief with regard to the proposition <God exists>) in order to greatly facilitate my pursuit of a variety of other things of final epistemic value (true belief and the avoidance of false belief with regard to the propositions that make up the subject matter of my scientific research). It is difficult to see how veritism -- at least as we have defined it so far -- can avoid saying that such a trade-off would maximize overall epistemic value So regardless of how we fill in the veritist's theory of

overall value and deontic theory in order to form a complete teleological epistemic theory, it appears that the veritist is committed to the highly implausible claim that my coming to believe in God's existence in order to get the grant is epistemically rational/reasonable/justified. (2013a, 364-365, emphasis added.)

. The clear implicature is that the veritist -- including, according to Berker, the process reliabilist -- is committed to a verdict about this case that is mistaken. In other words, the grant-seeking scientist case is a counterexample to veritism in general and process reliabilism in particular. Surprisingly, the following sentence appears immediately after the above-quoted passage: "My aim here is not to argue by counterexample" (365). This remark is very difficult to understand (or take literally) because it is hard to see how Berker's critique of epistemic teleology in this paper and its companion could proceed or succeed unless Berker's examples -- at least some of them -- are intended to be counterexamples. What other argumentative philosophical form could his two essays assume if not that of refutation (or challenge) by counterexample?² In any case, Berker proceeds to write as follows:

I want to use this example to illustrate a more general point ...: when determining the epistemic status of a belief in a given proposition, it is epistemically irrelevant whether or not that belief conduces (either directly or indirectly) toward the promotion of true belief and the avoidance of false belief in other propositions beyond the one in question. Veritistic epistemic teleology, in flouting this fact, ignores what we might call "the epistemic separateness of propositions." (2013a, 365)

According to Berker, veritistic epistemic consequentialism is committed to the verdict that the grant-seeking scientist's belief in God is justified. The belief is justified because

² In personal correspondence about the interpretation of this puzzling sentence, Berker writes that he is "happy to concede" that he does intend to be arguing against process reliabilism via counterexample. More precisely, however, he intended to say that his aim was not to argue by means of a single counterexample. But he also remarks, "Better yet, just ignore that sentence altogether." This last remark -- implying a commitment to argument by counterexample -- seems to set the record straight for both Berker's purposes and for mine.

it generates a greater balance of true beliefs over errors, most of the true beliefs being effects of the grant. Although the belief in God, we may suppose, is false, the numerous true scientific beliefs that are its products outweigh this false belief, making the belief a good choice from a "veritistic" point of view. Berker writes: "It is difficult to see how veritism, at least as far as we have defined it so far, can avoid saying that such a trade-off would maximize overall epistemic value." (Berker 2013b, 364)

3. WHY RELIABILISM IS INNOCENT OF ALL CHARGES

We see, then, where Berker is coming from, that is, what feature of reliabilism, veritism, or epistemic teleology supposedly gets it into trouble. But is it guilty of the charges leveled by Berker? First of all, is it correct that reliabilism would deem the grant-seeking scientist's belief to be justified? Not at all, I contend. I agree with Berker that classifying this or similarly generated beliefs as justified would be "beyond the pale" (2013a, 363). What I deny is Berker's claim that reliabilism is committed to such a verdict, or comes anywhere close to implying it.

With the help of Figure 1, let us review the core ideas of process reliabilism and explain why it would deliver the correct verdict, i.e., "unjustified." Simplifying somewhat, the core conditions of process reliabilism can be stated as follows (drawing on "What Is Justified Belief?," the original statement of process reliabilism).

S's belief that P (Bel(P)) is justified if and only if either

- (1) Bel(P) is the output of a belief-independent belief-forming process³ (BFP) that is unconditionally reliable, or
- (2) Bel(P) is the output of a belief-dependent BFP that is conditionally reliable, and all of the belief inputs to this BFP are themselves justified.⁴

³ A belief-independent belief-forming process is a process that takes no beliefs as inputs to its operation. A belief-dependent belief-forming process is one that takes beliefs (among other states, perhaps) as inputs.

⁴ A belief-dependent process is conditionally reliable iff a high ratio of its output beliefs are true when all of its input beliefs are true.

In the grant-seeking scientist case the target belief is a belief in God's existence. This belief is produced by a desire-driven process, something like wishful thinking. A desire or hope to bring about many true scientific beliefs (of his own or other people) generates or outputs a belief in God's existence. Thus, the operative belief-forming process is an instance of the process type, wishful thinking, which is clearly an unreliable process. This much is clearly implied by the example, though not explicitly stated in the description of the case. Hence, not only is the scientist's belief intuitively unjustified, but process reliabilism would definitely so classify it. End of story. The fact that the target belief token, Bel(P), also leads to many subsequent true beliefs has no bearing on its justificational status (J-status) according to process reliabilism.

[Insert Figure 1 about here]

As can readily be inferred from what we have said, process reliabilism is a backward-looking theory, not a forward-looking one. It is backward in the sense that the J-statuses assigned by this theory are always a function of events upstream from the target belief, not downstream from it. This is contrary to what Berker claims about consequentialist theories in general and process reliabilism in particular. Process reliabilism is a historical theory (Goldman, 1979; 1986, pp. 78, 83; 1999; 2009). In other words, it makes the J-status of a belief depend on the belief's provenance, not its subsequent effects. Thus, the J-status of the grant-seeking scientist's theistic belief depends on the belief-forming process that immediately outputs it (as shown in Figure 1 by the arrow immediately to the left of the oval shape labeled "Bel(God exists)," and perhaps by the desire itself. Nothing else depicted in the diagram is relevant to this belief's J-status. (Had other beliefs been causes of the target one -- for example, by being early links in a long chain of inferences -- they too could have been justificationaly relevant; but there are no such beliefs in this scenario.)

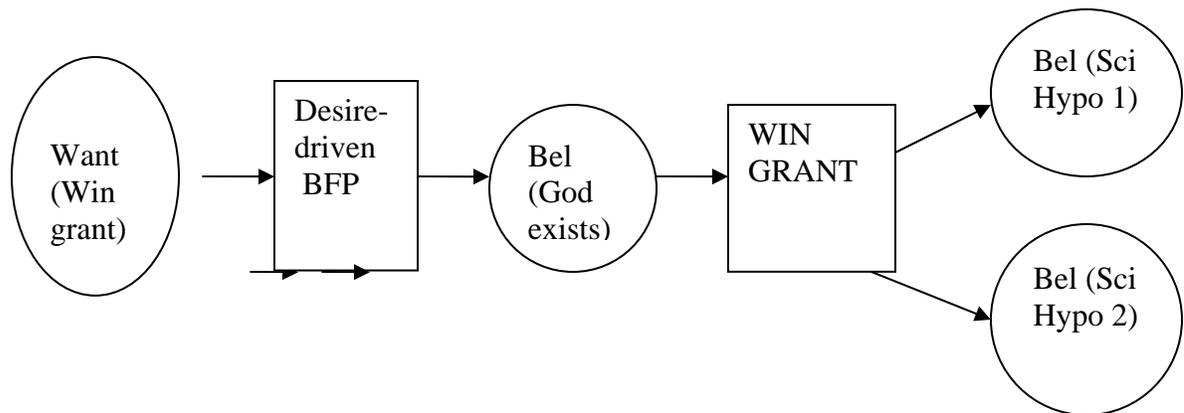
If the foregoing is correct, Berker's interpretation of process reliabilism is very wide of the mark. How might he have gotten the process reliabilist story so wrong? One possible problem is an early passage he cites (but misinterprets) from "What Is Justified Belief?" He writes:

In his reliabilist manifesto, 'What Is Justified Belief?' Alvin Goldman explicitly took the act-utilitarian analysis of RIGHT ACTION to be the model for his process-reliabilist analysis of JUSTIFIED BELIEF." (2013b, 364; emphasis added).

This, however, is an inaccurate characterization of what I was saying in the relevant passage. The referenced passage concerned a meta-level question of how exactly theories of justification should be formulated. Should they be formulated using other normative terms or only substantive (factual) terms, as would be suitable for an account that aims to be a "reductive" analysis of its analysandum? After posing this question, I announced my own intention to take the latter route, and I compared this course to act utilitarianism in normative ethics. Under act utilitarianism, only such descriptive terms as happiness, unhappiness, and causality play central roles. No normative terms (especially no morally normative terms) are employed. A parallel strategy, I urged, should be adopted in constructing a theory of justifiedness.⁵ It was only this "formal," meta-theoretic respect of similarity between reliabilism and act utilitarianism that I meant to endorse. Nothing I said there implied that process reliabilism is a species of consequentialism or instrumentalism in which beliefs downstream from a target belief (token) might be relevant to the latter's J-status.⁶

⁵ More precisely, I suggested that no normative epistemic terms should appear in any base clause (of a possibly recursive account of) a theory of justifiedness.

⁶ Berker (2013, 364) cites Roderick Firth (1981, 264) as characterizing reliabilism as a form of epistemological rule-utilitarianism, and indicates that he, Berker, thinks that Firth was exactly right in this characterization. But in the above-quoted passage Berker endorses the putative similarity of reliabilism to act utilitarianism. This difference may be important. Second, one cannot assume that Firth had any special authority as an interpreter of process reliabilism. So Berker's allusion to Firth as concurring (somewhat) with his interpretation does not go far in supporting it.



A separate point is worth making at this juncture, a point concerning the alleged principle that one should not "flout the separateness of propositions." The implication seems to be that each belief token should be epistemically evaluated "on its own merits." Its J-status should not be beholden to any other beliefs or their J-statuses. (This is a more definite statement of what a no-flouting principle would hold than any that Berker offers; but it seems to be in his spirit.) But such a principle would be misguided. When one infers a given proposition from other propositions that one antecedently believes, then the J-status of one's conclusion belief does depend upon -- or is beholden to -- the J-statuses of one's antecedent premise beliefs. In this very straightforward sense, the epistemic statuses of propositions are not always "separate from" -- i.e., independent of -- the epistemic statuses of other propositions. Process reliabilism is right on track in incorporating this sort of epistemological connectivity, and is in no way imperiled by any alleged principle of the separateness of propositions; because no legitimate principle of this sort has been demonstrated.⁷

4. RELIABILISM AND FORWARD- VS. BACKWARD-LOOKINGNESS

However, it must be acknowledged that there are treatments of reliabilism, including some of my own treatments, that support a certain chunk of what Berker says in drawing parallels between reliabilism and rule utilitarianism. Indeed, I drew such a parallel myself in *Epistemology and Cognition* (Goldman, 1986). There I advanced a rule consequentialist framework for process reliabilism, and classified it as a form of "verific consequentialism" (1986, 103). So Berker's interpretation is not wholly ungrounded. But does this framework imply that ascriptions of justificational status to (token) beliefs have a "forward-looking" as opposed to a "backward-looking" character, as Berker says? Not so, I reply.

⁷ Berker does countenance legitimate trade-offs between persons in some ethical cases (cf. Berker 2013a, 357); so perhaps he may not intend any universal principle against flouting the separateness of propositions in the epistemic domain. But if so, he is not explicit about this. And it would be very curious given that the "flouting" of the separateness of propositions is ostensibly introduced as a "no-no" and occurs in one of his titles (Berker 2013a). A ban against such flouting appears to occupy a central position in his epistemological world-view.

First of all, despite my rule framework, I still introduced the same sort of historical theory that was endorsed in "What Is Justified Belief?" In presenting the rule framework, I write as follows:

A cognizer's belief in *p* at time *t* is justified if and only if it is the final member of a finite sequence of doxastic states of the cognizer such that some (single) right J-rule system licenses the transition of each member of the sequence from some earlier states. (1986, 83)

Clearly, according to this principle, it is states and state-transitions leading to the target belief that confer justification on that belief, not what follows (temporally) from that belief. Moreover, two pages later, I make it clear that even within a rule-framework approach belief-forming processes are to be viewed as critical:

.... [W]e must strengthen our conception of J-rules in order to arrive at a satisfactory criterion of rightness. Instead of conceiving of them as specifying mere cognitive-state transitions, we must conceive of them as specifying cognitive processes, where by 'process' we mean a determinate kind of causal chain. (1986, 85; italics in the original)

With the process element re-introduced into the picture, the old (1979) story about the determinants of justifiedness is firmly back in place. The J-status of a given belief depends on what happens in the "run-up" to the belief, not what happens afterwards. And once again, crucial to this run-up are the psychological processes that generate the belief in question. So even if one chooses Epistemology and Cognition as the preferred formulation of reliabilism, one still finds no story to the effect that the J-status of a belief is determined by that belief's consequences ("down the road"). Yet this is what Berker claims in his interpretation of reliabilism, notably in the grant-seeking scientist case. The "central slogan" of epistemic consequentialism, he writes, is that "Epistemic normativity is a matter of bringing about states of affairs with a distinctive sort of value" (Berker 2013b, 377). As we see, however, process reliabilism does not hold that the

determination of a belief's J-status is a matter of what it brings about -- certainly it isn't directly a matter of such instrumental relations.

Berker is correct, however, in stressing the fact that under process reliabilism justificational normativity derives from the consequences of belief-forming process types, or what their outputs are.⁸ But is this reliabilist claim about justificational sources mistaken? Recall here some of the principal points made in the original proposal of process reliabilism (Goldman 1979). It was first pointed out that the cause(s) of a token belief is critical to its justificational status. Examples were given in which beliefs satisfied some orthodox account of justifiedness -- for example, their being well-supported by the agent's evidence -- but were intuitively unjustified because they were caused in strange, quirky, or unacceptable fashions. For example, they were caused by wishful thinking, or the accidental movement of a brain surgeon's hand, or the blinding aura of the Presidency. The argument continued as follows:

In general, a strategy for defeating a noncausal principle of justifiedness is to find a case in which the principle's antecedent is satisfied but the belief is caused by some faulty belief-forming process. The faultiness of the belief-forming process will incline us, intuitively, to regard the belief as unjustified. (Goldman 1979, 2012, 36, using the pagination of Goldman 2012).

So far, however, this only shows that a token belief's provenance is a key to its J-status. But what features of its provenance are crucial? Reflection suggests that the belief-producing process must be one that, generally speaking (not necessarily on the occasion in question), is prone to be reliable, i.e., tends to output many more true beliefs than false beliefs.

Granted that principles of justified belief must make references to causes of belief, what kinds of causes confer justifiedness? We can gain insight into this problem by reviewing some faulty processes of belief-formation, i.e., processes

⁸ This is my (re-)formulation of Berker's statement, "Epistemic normativity is a matter of bringing about states of affairs with a distinctive sort of value."

whose belief-outputs would be classed as unjustified. Here are some examples: confused reasoning, wishful thinking, reliance on emotional attachments, mere hunch or guesswork, and hasty generalization. What do these faulty processes have in common? They share the feature of unreliability: they tend to produce error a large proportion of the time. By contrast, which species of belief-forming (or belief-sustaining) processes are intuitively justification-conferring? They include standard perceptual processes, remembering, good reasoning, and introspection. What these processes seem to have in common is reliability: the beliefs they produce are generally true.... Furthermore, our intuitions of comparative justifiedness go along with our beliefs about the comparative reliability of the belief-causing processes. (Goldman 1979/2012, 37)

Thus, process reliability seems exactly on track in pinpointing the propensity of a process type to output more true beliefs than false beliefs as the key to the determination of justificational statuses. (It is not all downstream effects of a process type that figure in the formula; only its immediate outputs.) Although this formula does not use the terminology of epistemic "value," it is no stretch to connect up with that language. The important point is that these ideas seem right, not mistaken, as Berker claims. This is the best explanation of why process reliabilism rapidly became a serious contender among theories of justification, soon after "What Is Justified Belief?" was published. It offers a parsimonious account of the intuitions about cases reviewed above.

What we have acknowledged is a covert forward-lookingness embedded in the background of process reliabilism. Types of processes have justification-conferring or unjustification-conferring properties as a function of their truth ratios, where truth-ratios are a matter of the pattern of outputs that the different process types have. Such forward-lookingness is relevant, however, only for process types, not to the (target) belief tokens themselves. Belief tokens are what get evaluated as justified or unjustified, reasonable or unreasonable, knowledge-worthy or not knowledge-worthy. These bearers of J-status do not enjoy their respective J-statuses as a function of their consequences, contrary to what Berker seems to imply in his discussion of the grant-seeking scientist.

Berker has own way of distinguishing different components of the stories offered by process reliabilists and other so-called epistemic teleologists. He distinguishes between different "deontic focal points" such as beliefs, on the one hand, and belief-forming processes on the other. However, despite acknowledging such differences, he makes inadequate use of them by my lights.

5. OTHER POSSIBLE COUNTEREXAMPLES

It is time to examine two other counterexamples discussed by Berker: Roderick Firth's (1981) John Doe case and Berker's own prime number case. Here is the John Doe case:

John Doe is a brilliant set theory [sic.] who is on the cusp of proving the Continuum Hypothesis: all he needs is six more months. But, alas, poor John is suffering from a serious illness that, according to his doctors, will almost certainly kill him in two months' time. John stubbornly clings to a belief that he will recover from his illness, and not only does this belief comfort him, but -- let us suppose -- it in fact significantly raises the chances that he will live for the six months that he needs both to complete his proof and to derive from it a variety of consequences for the rest of set theory. In other words, John's belief that he will recover is a causal means to his procuring a large number of true set-theoretic beliefs sometime in the future. But is John's belief epistemically justified? Is it the kind of belief that, from a purely epistemic perspective, he should be holding? (Berker 2013b, 369)

Berker says that Firth's John Doe example and the others like it "point the way to a more serious problem that afflicts almost all varieties of consequentialism with respect to epistemic matters" (2013b, 369)

Suppose our epistemic consequentialist accepts the theory of final value dubbed "veritism" by Alvin Goldman: namely, the view that (i) true beliefs, and only true

beliefs, have final epistemic value, and (ii) false beliefs, and only false beliefs, have final epistemic disvalue (Goldman 1999, 5; 2001, 31). Suppose, moreover, that our epistemic consequentialist endorses a theory of overall value which allows all causal means to be ways of promoting final epistemic value. Then regardless of the other details of the epistemic consequentialist's theory of overall value, and regardless of her deontic theory, it is difficult to see how our epistemic consequentialist can resist the conclusion that John Doe's belief that he will recover is epistemically justified. (2013b, 369-370).

This passage runs together two things that should be kept apart: (A) a theory of epistemic value called "veritism" and (B) a theory of belief justifiedness called "process reliabilism." It is correct that in Knowledge in a Social World (Goldman 1999) I advocated a form of epistemic consequentialism called "veritism". And it is correct that this epistemic consequentialist scheme was used to evaluate (epistemically) a wide variety of social practices such as interpersonal argumentation and judicial rules for the admitting or excluding evidence in a trial. Such rules and practices, it was argued, can be assessed in veritistic terms, and this is a reasonable thing for theorists to do, since patterns of truth-value attainment are of central importance in many social (as well as individual) domains. But Knowledge in a Social World never proposed that the indicated veritistic evaluations are justification-directed evaluations. Specifically, the justification of beliefs simply wasn't on the table. Veritistic evaluation, as understood there, was indeed a form of direct consequentialist assessment. Practices, rules, and institutional arrangements, it was urged, should be evaluated in terms of "downstream" effects such as the truth-ratio of jury verdicts. But it was nowhere suggested that belief justifiedness was to be so evaluated.

Contrary to Berker, therefore, it is not really "difficult" to see how a veritist can avoid the (unwanted) conclusion that John Doe's recovery belief is epistemically justified. As we saw earlier, process reliabilism asserts that it is the causes, not the effects, of a belief that are relevant to its J-status. So it should be clear that if John Doe's recovery belief is to be evaluated justificationaly, it must be evaluated in terms its causes, not its downstream effects. Veritism of the sort advocated in Knowledge in a

Social World simply has nothing to say about (token) belief evaluation (in justificational terms). Hence it seems straightforward for veritism to avoid taking a position about the John Doe's belief. The veritistic program advanced in Knowledge in a Social World isn't in the business of making justification evaluations.

Despite the passage quoted above, Berker's subsequent discussion of the John Doe case indicates that he does not, in the end, intend his original John Doe case to be a counterexample to process reliabilism. What he does claim, however, is that just a little tweaking of the case will yield such a counterexample:

... [If] we tweak our story a bit, the same result holds even if our epistemic consequentialist embraces an indirect deontic theory: all we need to do is make situations like the one John [Doe] faces sufficiently common in his world that a general practice of stubbornly believing that one will recover from an illness tends, on average, to greatly promote the acquisition of true belief. (Berker 2013b, 370)

The reasoning here is fairly straightforward, as I would reconstruct it. According to process reliabilism, a process is justification-conferring just in case it has a high truth-ratio. So a process of "wishful recovery belief-formation" is justification-conferring just in case it has a high truth ratio. By tweaking the original John Doe example so that not Doe's wishful recovery belief-forming process not only generates many true beliefs of his own, but this same pattern recurs frequently in Doe's world. Hence the process type is reliable (Berker implies) and process reliabilism must declare Doe's output belief(s) as justified.

Berker thus seems to assume that the domain in which a process type should be evaluated for reliability is the actual world of the epistemic subject, in this case, John Doe. Truth ratio in the actual world as a whole, of course, comprehends far more than the truth ratio of a single subject's beliefs. Berker's assumption, however, does not comport well with the unresolved, continuing debate on this issue among reliabilists themselves. My own publications (including Goldman 1979, 1986, 2012) have presented several alternative possibilities. These include the following three alternatives.

First, the domain of evaluation might be expanded beyond the actual world of the agent to other nearby possible worlds. Second, the domain of evaluation might be neither the actual world nor neighboring worlds but the "normal world," where "normality" is specified by what ordinary people presume to be true of the actual world (cf. Goldman 1986, 108). Third, reliabilism about justification might be revised more extensively by introducing an interpretation I originally called "two-stage" reliabilism (Goldman 1992), but am currently inclined (following a suggestion of Elizabeth Fricker) to label "approved-list" reliabilism.

The idea behind the third of these options is that people pass through two stages in acquiring and deploying the cognitive tools for making justification evaluations. In the first stage they acquire beliefs about which belief-forming process types are reliable or unreliable. This is based on their own and/or other people's "observations" of which belief-forming processes have high or low truth-ratios when used in the actual world. Second, after acquired and mentally storing lists of such "approved" and "unapproved" processes, they proceed to apply their favored and disfavored process types to scenarios in which rather different truth-ratios for processes may be operative (e.g., in worlds ruled by an evil demon). When making judgments at the evaluation stage, they do not judge a process's reliability by what is stipulated to occur in the hypothetical scenario, but rather by the attributors' previously entrenched lists of approved and unapproved process types. Thus, in the case of John Doe, we evaluate his recovery belief as unjustified because we previously stored wishful thinking in our minds as an unreliable process. This unapproved (indeed, disapproved) process will continue to exert its influence on our evaluation of Doe's belief even if we are told, as in Berker's tweaking, that recovery beliefs are generally quite reliable in the actual world Doe inhabits.

My present inclination is to favor the approved-list form of reliabilism (it will be applied to the next putative counterexample too). But it isn't necessary at this juncture to register a decided preference for any one of these alternative forms of reliabilism. The main point is that it isn't easy to refute reliabilism by a simple example such as the "tweaked" version of the John Doe story, because reliabilism has several options for ironing out this wrinkle. Berker needs to give us a broader argument to demonstrate reliabilism's alleged futility. Perhaps such a broader argument is what he intended with

his appeal to the separateness of propositions. Unfortunately, for reasons provided earlier, that "broad" line of argument fails. Here I am saying that his critique by counterexample seems at best inconclusive. And to the extent that many others have offered similar types of counterexamples, it is also not very novel.

There remains a final counterexample that Berker lodges against process reliabilism, a more novel and interesting one, in my opinion. It is his prime number case, which runs as follows. Whenever S considers whether a particular natural number n is prime, S forms a belief that n isn't prime. This process tends toward a truth-ratio of 1.0 as n increases, so process reliabilism, says Berker, must declare that all of these beliefs are justified. Intuitively, however, they are unjustified. How might reliabilism reply to this apparent counterexample? There are several possible responses.

First, suppose that the domain we select for determining a process-type's reliability is the actual world (plus, perhaps, a class of a relevantly nearby possible world). I now make the following observation. In all likelihood, people who use this process type apply it exclusively or primarily to relatively small integers. Within that class of numbers, however, the truth-ratio of the process will not be very high, because many small integers are prime (1, 2, 3, 5, 7, 11, etc.). So it is questionable whether the counterexample even gets off the ground.

A second reply would invoke a rarely discussed constraint I proposed on belief-forming processes in "What Is Justified Belief?", a constraint requiring eligible process types not to be content-specific. Instead, they must be content-neutral. For epistemological purposes, suitable process types should be general processes, which can presumably take many different contents as inputs (and as outputs). No process should be confined to a single, specific subject-matter. Thus, what Berker calls a "process" in his prime number example cannot really be accepted by process reliabilism. Agent S may be said to use a certain rule for deciding whether a particular natural number is prime (always saying "no" for every natural number specified). But this rule should not be equated with a process. Processes will presumably be used, of course, in following the rule, but that does not make the rule itself a process. A distinction between "processes" and "methods" was drawn in Epistemology and Cognition (Goldman 1986, 93-95), proposing that processes be understood as basic psychological processes and

methods are understood as learnable algorithms, heuristics, or methodologies (not part of one's native cognitive equipment). The rule in Berker's prime number example clearly falls on the methods side of this distinction. So it is reasonable to impose a criterion of justifiedness for methods that is suitable to methods rather than processes. Use of a reliable method does not automatically or invariably confer justifiedness on an output belief, because a perfectly reliable method might have been acquired in a defective -- e.g., haphazard though lucky -- fashion. If that is how an agent acquired a given method, its reliability per se does not guarantee justifiedness of a belief so generated. Unless an agent acquires Berker's prime-number rule appropriately, a belief generated by it won't be justified. If it is acquired appropriately, though, its output beliefs would be justified. Hence, no counterexample would emerge.

Finally, the approved-list version of reliabilism looks like a very promising way to handle the counterexample -- even if we do not press the distinction between processes and methods. It is part of the approved-list version of reliabilism (see Goldman 1992) that when a novel process type fails to match any type already on one's approved list or unapproved list, then an epistemic evaluator will find the most similar type on one of the lists and classify the target process type accordingly. A likely point of comparison for purposes of Berker's prime number example is some process type that is entirely arbitrary and unmotivated. That's the character of the prime number rule as Berker presents it. But arbitrary and unmotivated process types can be expected to belong on an agent's unapproved list, not on her approved list. So Berker's rule would be regarded as unreliable and hence incapable of generating justified beliefs -- the correct verdict by everybody's lights.

In sum, once Berker's misinterpretation of process reliabilism is straightened out, his putative counterexamples turn out to be either totally unproblematic for reliabilism or at best questionable cases that reliabilism has pretty good prospects for accommodating.

- Berker, Selim (2013a). "Epistemic Teleology and the Separateness of Propositions," Philosophical Review 122(3): 337-394.
- Berker, Selim (2013b). "The Rejection of Epistemic Consequentialism," Philosophical Issues 23: 363-387.
- Firth, Roderick (1981). "Epistemic Merit, Intrinsic and Instrumental," Proceedings and Addresses of the American Philosophical Association 55: 5-23.
- Foot, Philippa (1985). "Utilitarianism and the Virtues," Mind 94: 196-209.
- Fumerton, Richard (2001). "Epistemic Justification and Normativity," in M. Steup, ed., Knowledge, Truth, and Duty (49-60). Oxford: Oxford University Press.
- Goldman, Alvin I. (1979). "What Is Justified Belief?" In G. Pappas (ed.), Justification and Knowledge. Dordrecht: Kluwer. Reprinted in Goldman (2012). Page references to this work are to the 2012 reprinting.
- Goldman, Alvin I. (1986). Epistemology and Cognition. Cambridge, MA: Harvard University Press.
- Goldman, Alvin I. (1992). "Epistemic Folkways and Scientific Epistemology," in Liaisons: Philosophy Meets the Cognitive and Social Sciences. Cambridge, MA (155-175). MIT Press.
- Goldman, Alvin I. (1999). Knowledge in a Social World. Oxford: Oxford University Press.
- Goldman, Alvin I. (2001). "The Unity of the Epistemic Virtues," in Abrol Fairweather and Linda Zagzebski, eds., Virtue Epistemology: Essays on Epistemic Virtue and Responsibility, 30-48. Oxford: Oxford University Press.
- Goldman, Alvin I. (2012). Reliabilism and Contemporary Epistemology. New York: Oxford University Press.
- Korsgaard, Christine (1996). Creating the Kingdom of Ends. Cambridge: Cambridge University Press.

Rawls, John (1971). A Theory of Justice. Cambridge, MA: Harvard University Press.

Thomson, Judith Jarvis (1976). "Killing, Letting Die, and the Trolley Problem." The Monist 59: 204-217.