

probably assume that this token process acquires *its* value by standing in a singular causal relation to the true belief in question. This is what I would call *token instrumentalism*. The puzzle about token instrumentalism is how the process token can “return” any of the value it inherits from the true belief it causes so that the latter’s total value exceeds the value it would have were it instead caused by an unreliable process.

However, a token reliable process might acquire “derivative” epistemic value in an entirely different way without encountering the same problem. It is widely assumed in moral theory that something’s being a token of a valuable type automatically confers value (at least *prima facie* value) on that token. For example, an action’s being a case of promise keeping automatically confers (*prima facie*) value on that action token. Similarly, if we assume that being reliable makes some *type* of belief-forming process epistemically valuable (for example, taking visual appearances at face value), then any *token* of such a belief-forming type would also have epistemic value. Thus, to qualify as epistemically valuable, a token belief-forming process would not have to succeed in generating any true belief. It would not have to stand in a singular causal relation to a true belief. Value might accrue to it in a different, two-step fashion. In step 1, the type acquires value because many instances of it produce token true beliefs. In step 2, a token of the process type acquires epistemic value because it instantiates that valuable type, whether or not it produces a true belief of its own. If all this is correct, then it’s false that the only way for a belief-forming process-token to acquire (instrumental) epistemic value is for *it* to cause an epistemically valuable upshot, for example, a true belief.

This point helps position reliabilism for an escape from the swamping problem. If the source of a process token’s value needn’t be a true belief that it produces, then there’s no guarantee that the value of the resulting true belief must swamp, or trump, that of the token. There remains, however, a nagging problem. How can the value of the token increase the value of the true belief? Value theory routinely assumes that value transmission runs “from” achieved goals, outcomes, or ends “to” the means that produce them. How could it countenance the idea that value transmission runs in the opposite direction, viz., from means to ends? That is what would be required if a (valuable) reliable process is to *increase* the value of a true belief by virtue of causing it. Is this scenario the slightest bit plausible?

I think it is. Here is an example to illustrate the indicated pattern of transmission. Any work by Rembrandt, whatever its intrinsic aesthetic quality, will be assigned a high degree of value simply because it was produced *by Rembrandt*. The fact that Rembrandt painted it is enough to confer significant value on it, as its price in the art market would reflect. The extra value assigned to a genuine Rembrandt need not be merely a recognition of market value. Instead, it may be a recognition of what explains that market value, i.e., the fact that people admire or esteem such a work because it’s the product of a masterful hand (and eye).

Here, then, is an example of a process token that acquires value by being a token of a valuable type—being a case of generation—being a case of promise keeping or esteem for paintings (or Robert’s reliable process that produces true belief). A highly prized process—being a case of promise keeping or esteem for paintings (or Robert’s reliable process that produces true belief)—elicits enhanced admiration or esteem for paintings (or Robert’s reliable process that produces true belief) in the case of Rembrandt’s paintings (or Robert’s reliable process that produces true belief). The value of the true belief itself, and the value of the process that produces it. Since the value of the process that produces it, the reliably produced true belief could easily exceed the value of a similar true belief produced by an unreliable process. This easily meets the challenge of the swamping problem.

I now digress a bit by addressing another challenge to reliabilism based on a theory of value, though this problem makes no appearance in any of the papers of the present volume. It is fashionable nowadays to seek ways to bridge or unify the normative fields of ethics and epistemology. One proposed unification compares epistemic teleology (or consequentialism) to moral teleology (or consequentialism). Reliabilism is often said to be a consequentialist approach. It posits, as core epistemic values (final values), states of true belief and error avoidance. Justifiedness of belief is some sort of deontic status, such as being permitted, or entitled, to believe. This kind of deontic status is ultimately rationalized, or derived, from the conduciveness of a belief—or, rather, conduciveness of the belief-forming process—to the production of good epistemic consequences, in the form of true beliefs or the absence of false beliefs. Viewing reliabilism in this light is encouraged by formulations in terms of a right-rule architecture, as is offered in essay 4 (and in *Epistemology and Cognition*). As we shall now see, however, this teleological or consequentialist rendering of reliabilism is used by some critics as a basis for putative counterexamples to it.

Here is an example found, in various permutations, in Firth (1981), Fumerton (2001), Berker (unpublished), and other critics of reliabilism. A scientist seeks a grant from a religious organization. Although she is an atheist, she realizes that her only chance of getting funding from the organization is to form a belief in God’s existence. She knows that if she received the grant, she would use it to further her research, and this would allow her to form many new true beliefs and revise many previously held false beliefs on matters of intellectual significance. Would she be justified, in these circumstances, in forming a belief that God exists? Reliabilism, according to these interpreters, implies that she would be so justified, because so many true beliefs would result, along with the abandonment of many false beliefs. All of these consequences would yield great epistemic value, so epistemic teleology would certainly permit it to be done and perhaps even require it. Reliabilism must deem it to be justified. Intuitively, however, it isn’t at all justified.

This, however, is an inaccurate portrait of reliabilism and its implications, at least on my rendering of reliabilism. As I have presented it, process reliabilism

(PR) does not warrant the specified conclusion in the scientist case. There are three ways the alleged interpretation fails to comport with PR. First, the scientist example proceeds on the assumption that reliabilism is a value-maximizing kind of theory. But PR does not equate justified belief with belief that maximizes (final) epistemic value. This was clear in *Epistemology and Cognition*, where right J-rules—conformity with which yields justifiedness—are ones that authorize processes with a high truth *ratio*, not necessarily a high quantity of true belief (Goldman, 1986: 106). The high truth-ratios standard gives more weight to error avoidance than to truth acquisition. Belief-forming processes can earn high truth ratios by forming beliefs very cautiously and carefully—generating very few beliefs, hence rarely falling into error, but also rarely generating truths. Thus, unlike the most famous forms of moral consequentialism, e.g., utilitarianism, it is not a value-maximizing theory.

Second, one could only impute to reliabilism the implication that the scientist's belief is justified in the specified example by ignoring the *process* component of reliabilism. PR does not determine a belief's justificational status simply by its direct or indirect causal consequences (actual or hypothetical). In fact, it doesn't look at the consequences of the target belief at all. Instead it looks—most immediately—at the process (or processes) of which the target belief is an *effect*. If the process that causes the belief is a token of a type such that *its* (the *type's*) belief-upshots have a high truth ratio, then the belief is justified. If not, not. What is the belief-forming process in the scientist case? Since the presented case is one of propositional justification, the case is one where the agent has not (yet) formed a belief. Therefore, some assumption must be made about *how*—i.e., by what psychological route—she would form the belief if she adopted it. What might such a psychological route be? Presumably, something like “being driven to believe a proposition by irrelevant goals or desires”—irrelevant, that is, to the truth of the believed proposition. No such candidate process would have a very high truth ratio. So PR implies that the scientist would be unjustified in forming a belief in God's existence. That such a belief would nonetheless have favorable downstream epistemic consequences has no bearing, under PR, on its justificational status.<sup>15</sup> Only upstream cognitive activity—i.e., psychological activity actually or potentially involved in the target belief's production—is relevant to that belief's justificational status. Notice that the anticipated downstream consequences (true beliefs in scientific propositions) would presumably be formed by very different belief-forming processes than the one that would generate a belief in God. So, even if those subsequent, science-generated beliefs would themselves be justified, this in no way suggests that the scientist's belief in God (were she to form this belief in the manner specified) would be justified.

Berker stresses that justificational assessments of belief must observe the principle of “the epistemic separateness of propositions.” This principle states that 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” (unpublished, p. 29). Moreover, he contends that veritistic epistemic teleology flouts this principle and that reliabilism is a species of veritistic epistemic teleology. These contentions are mistaken. As explained previously, PR does not flout the principle, because it is not really a species of veritistic epistemic teleology.

I turn next to a closer examination of Williamson's (2000) theories of knowledge, because Williamson's global view is reliabilist through and through. His views are very congenial, at least in general contours, to the views I have defended, for example, in the specific details of the theory are less successful than mine. I also suggest that Williamson is downplaying the process element in PR. In response, he has replied that he did not intend to distance himself from the process element. If so, all well and good, but this fact was not readily apparent.

The final two sections of essay 7 address the topic of evidence. Williamson holds that a person's body of evidence is all and only what she knows ( $E = K$ ). A battery of intriguing and powerful arguments were offered in *Knowledge and Its Limits* to show that no rival view could account for the intuitive data better than (the quite surprising and initially counterintuitive)  $E = K$  thesis. I respond by arguing that at least one rival view performs at least as well as  $E = K$  in accounting for Williamson's “data,” namely, the view that evidence is non-inferential propositional justification ( $E = NPJ$ ). This dispute is relevant to Williamson's knowledge-first treatment of epistemology, of which I am not yet persuaded. In his response, Williamson (2009) repeats his claim that evidence is factive, and that this feature favors a knowledge account as opposed to the  $E = NPJ$  view. He also recognizes the possibility of strengthening  $E = NPJ$  by adding the factivity component. He doesn't say much to favor his preferred view over this alternative. He only remarks that  $E = TNPJ$  is a “rather unnatural hybrid” (2009: 311). He doesn't spell out, however, what makes this particular hybrid “unnatural” or what makes  $E = K$  natural.<sup>16</sup>

#### Cluster 4: Essays 8, 9, and 10

The three essays in this cluster belong in the category of social epistemology, a subfield of epistemology to which I have recently devoted a lot of attention (more than the present volume indicates). Essay 8 addresses two rather distinct topics in social epistemology, peer disagreement and epistemological relativism. Here I confine my summary remarks to relativism.