

VOGEL ON RELIABILISM

SHLOKA V. JANAPATY

October 19, 2024

Epistemological reliabilism holds that knowledge is the fruit of a process that consistently yields true beliefs. In *Reliabilism Leveled*, Jonathan Vogel makes the exposition and critique of the reliabilist position his chief philosophical project. I will synthesize several arguments from his work to argue that both versions of reliabilism are deeply flawed: counterfactual reliabilism presents a far too demanding account of knowledge, and neighborhood reliabilism appears vulnerable to a bootstrapping problem. In response, I will offer what I believe to be the strongest reliabilist defenses and evaluate their adequacy.

A good reliabilist places truth at the normative center of her epistemic process. She would like her belief $B(X)$ to be infallibly true, but she cannot guarantee that $B(X)$ implies X and not X implies not $B(X)$. Instead, the reliabilist sets the bar slightly lower, requiring that the process by which she acquires $B(X)$ be *reliable*. For example, the rising of the sun is a reliable process. We can expect to see a sunrise tomorrow, barring exceptional circumstances (e.g. unknowingly being transported to a place that experiences polar night, a cataclysmic volcanic winter, etc.). The reliabilist hopes that by requiring reliability, she will minimize her vulnerability to false belief and adhere to generally consistent *modus operandi*.

Let us now consider the possible worlds account. Here, a belief $B(X)$ is reliable if it holds true in the neighborhood of worlds near the actual world, a process we will call P . This yields what we will henceforth refer to as **Condition 1**: $B(X)$ by process P implies X . Vogel adds that many versions of reliabilism also include a tracking condition—not X implies not $B(X)$ —henceforth **Condition 2**. Two versions of reliabilism emerge from these conditions for knowledge. **Neighborhood reliabilism** (NR) is the view that knowledge of X , or $K(X)$, is contingent on satisfying Condition 1, while **counterfactual reliabilism** (CR) imposes both Conditions 1 and 2. It is worth noting that in both formulations, reliabilists reject the justification condition partly as a means of blocking the skeptical hypothesis. They also exclude higher-level knowledge requirements, such as $K(X)$ implies $K(K(X))$ and $K(X)$ implies $K(P \text{ is reliable})$. This is because $K(X)$ is regarded as a fact entirely independent of $K(K(X))$ and knowledge of the process's reliability. One can know X without knowing that they know X or that they arrived at X through a reliable process.

And yet, there are many instances when we know that we know. Further, we may know that we arrived at this knowledge via a reliable process. I know that

my great-grandmother passed away in a hospital in Bangalore, India. If someone were to question this, I might retort, “I know that I know! My grandmother was in the room when it happened. She would never lie to me.” The doctor in attendance also knows this fact. Moreover, he knows that he arrived at this conclusion via a reliable process—the heart rate monitor had been accurate and precise for say, the last 100 patients. To expose the difficulty with counterfactual reliabilism, Jonathan Vogel offers the following example. Suppose you ask your friend Omar, an honest individual, if he is wearing new shoes. Upon receiving his confirmation, your belief that Omar has a new pair of shoes becomes true. Because you trust Omar, your belief that he has new shoes is reliable. Thus, you know Omar has new shoes, and you also know that you know that Omar has new shoes. This case passes Condition 1; your belief that Omar has new shoes is true by a reliable process, implying knowledge. However, it fails Condition 2. Even if Omar did not have new shoes, you would still believe that he did. Vogel points out that it is hard to fathom believing something to be false that you actually believe is true, especially if you came by that belief through an ostensibly reliable process. By CR, you do not know that Omar has new shoes. But in the everyday sense of knowledge, you do. We conclude that CR presents a far too demanding account, one that excludes versions of knowledge that we do have.

Let us now consider neighborhood reliabilism, a weaker form of CR. Jonathan Vogel argues that NR is also beset with epistemic troubles, namely the problem of bootstrapping. Consider the following example. Suppose Roxanne checks her gas gauge often while driving. Over time, she accumulates a great deal of raw data about the state of the gas gauge, as well as inferences about how much gas is in the tank. In fact, she believes that the state of the gas gauge corresponds exactly to how much gas she has. At one point while driving, Roxanne constructs the belief that at this moment, the gas gauge reads empty and the tank is indeed empty. It follows by deduction that at this moment, the gauge is accurate. On the basis of multiple observations, Roxanne applies induction to the base case and concludes that the gauge is always accurate. What’s more, the gauge’s history of being accurate makes it reliable. Since she used a reliable process to conclude that the gauge is always accurate, we can say that Roxanne knows that her tank is empty right now.

Clearly, there is something terribly amiss about Roxanne’s process. She appears to have awarded the lofty status of knowledge to beliefs that she assumed to be true and sampled repeatedly—bootstrapping—to give the appearance of a reliable process. Indeed, Roxanne has fortified beliefs formed by reliable processes with the false knowledge that they were formed reliably, taking her conclusion to be the starting assumption. This circularity strips her epistemic process of its very reliability.

Thus far, I have argued that both versions of reliabilism are deeply flawed. In response, I will now offer two defenses of counterfactual reliabilism against bootstrapping. The first reliabilist response is that Roxanne’s epistemic process does not qualify as reliable. She assumed that the state of the gas gauge corresponds exactly to how much gas she has, verifying the conclusion that enables

her knowledge of the tank level as a starting assumption. In other words, we cannot conclude that she knows that her tank is empty because she used an unreliable, circular epistemic process. We can safely disqualify this as a case of knowledge.

Critics of reliabilism may employ two responses to defeat this counterargument. First, if the reliabilist hopes to make this claim, they must also propose action-guiding criteria for ascertaining reliability. What differentiates bootstrapping (claimed as unreliable) from other ostensibly reliable processes? Why is a dipstick more epistemically reliable than the gas gauge? Such questions may seem trivial in the gas gauge case, but they blow up when the reliabilist attempts to dismiss radical skepticism. Second, the reliabilist has herself rejected higher-level knowledge entailments, such as $K(X)$ implies $K(P \text{ is reliable})$; that is, we cannot say whether bootstrapping is reliable or not. Thus, she cannot claim that this is not a case of knowledge with delivering a crippling blow to her own stipulations.

The reliabilist has another arrow in her quiver. Perhaps bootstrapping is not all that bad. It certainly diminishes trust in our body of knowledge by introducing the possibility that we self-validated our preexisting intuitions. But these self-validations are premised on consistent, regular, and observable patterns—not knowledge, but a sense of intuition premised on past successes. Reliabilism may sanction silly, outlandish conclusions via bootstrapping, but it is not particularly pernicious in the epistemic regime of Roxanne’s gas gauge case. In fact, the reliabilist can rehabilitate her argument by countering that this case grossly plays up the perniciousness of bootstrapping, mischaracterizing the epistemic regimes in which bootstrapping is at play. A more relevant example related to radical skepticism is offered later in the article by Vogel. A common skeptical response goes as follows. I have a hand. I can perceive that I have a hand through touch, movement, and other sensations. If I have a hand, I am not a brain in a vat. Since perception is a reliable process, my belief that I am not a brain in a vat becomes knowledge. This is a clear case of bootstrapping; the argument proceeds circularly because it takes as a starting assumption that brains in vats do not perceive hands. If the reliabilist admits that such bootstrapped responses to the skeptic do not count as knowledge, she forgoes her defeat of the skeptic but retains reliability as a criterion as knowledge. In other words, by severing the arm she can successfully save the body.

Reliabilism posits that knowledge arises from processes that reliably produce true beliefs. Jonathan Vogel contends that both major versions of reliabilism have significant shortcomings. Counterfactual reliabilism sets an excessively high bar for defining knowledge, while neighborhood reliabilism seems susceptible to a problematic bootstrapping effect. In response, I argue that the reliabilist can deny the reliability of bootstrapping or give up their defeat of the skeptic to save their position.

References

- [1] Vogel, Jonathan (2000). *Reliabilism Leveled*. *Journal of Philosophy* 97 (11): 602.