

## EPISTEMIC BOOTSTRAPPING\*

I have argued that process reliabilist theories of knowledge are fundamentally incorrect.<sup>1</sup> This inadequacy is brought to the fore by the phenomenon of epistemic bootstrapping. Reliabilism would allow someone who forms a belief by a reliable process to know trivially, by bootstrapping, that her belief was formed by a reliable process.<sup>2</sup> But the knowledge in question is not at all trivial, and a person can certainly form a belief by a reliable process without knowing that she did so. Reliabilism errs by attributing knowledge to subjects who lack it. This problem is especially acute for the reliabilist, because a sharp distinction between *forming* a belief by a reliable process and *knowing* that one's belief is formed by a reliable process is a cornerstone of the reliabilist position. In my earlier discussion, I said that a justificationist approach to knowledge treats properly the kind of case which the reliabilist gets wrong. The reason why the subject in these cases fails to know is that he has no evidence available to him which would justify his belief. I still believe that this picture is correct, but there is more to be said about the relation between bootstrapping and reliabilism, the relation between bootstrapping and justification, and about the disutility of bootstrapping itself.

## I. A PROBLEM FOR RELIABILISM

To introduce the topic of bootstrapping, here is an example:

*Gas Gauge Case.* Roxanne drives a car with a well-functioning, reliable gas gauge. She has never looked into the status of the gauge or others like it; she has no information whatsoever on the subject. Rather, Roxanne

\*For much helpful discussion, I am grateful to Michael Bergmann, Gina Calderone, Stewart Cohen, David Christensen, Kenny Easwaran, Branden Fitelson, Michael Glanzberg, Alexander Jackson, Roberta Millstein, Bernard Molyneux, Brad Morris, Sherrilyn Roush, and Michael Teitelbaum. A version of this paper was presented at the University of Wisconsin, Madison, and I thank the audience for their valuable comments.

<sup>1</sup>Jonathan Vogel, "Reliabilism Leveled," this JOURNAL, xcvi, 11 (November 2000): 602–23.

Richard Fumerton devised a similar line of argument around the same time I did. See his *Metaepistemology and Skepticism* (Totowa, NJ: Rowman and Littlefield, 1995), pp. 177–80, and more recently Fumerton, "Epistemic Internalism, Philosophical Assurance, and the Skeptical Predicament," in Thomas M. Crisp, Matthew Davidson, and David Vander Laan, eds., *Knowledge and Reality: Essays in Honor of Alvin Plantinga* (Dordrecht, The Netherlands: Kluwer, 2006), pp. 179–91.

<sup>2</sup>When I use the term 'reliabilism', I mean process reliabilism, unless otherwise indicated. For some clarification about the nature of reliabilism, see below.

automatically forms beliefs about the level of gas in the car's tank simply by consulting the gauge. For example, if the gauge reads "F" she immediately and directly forms the belief that the car's tank is full. Given that the gauge is reliable, it seems clear that Roxanne's belief that the car's tank is full is formed by a reliable process. Now, Roxanne can also observe what the state of the gauge itself is, if she chooses to. Roxanne notes that the needle reads "F" at the time when she believes, by reading the gauge, that the tank is full. Roxanne conjoins her belief that the gauge reads "F" with her belief that the tank is full, and deduces that the gauge reads accurately on this occasion. We can suppose that Roxanne repeats this strange procedure a good number of times, accumulating beliefs that the gauge reads accurately at various times  $t_i$ . Roxanne goes on to conclude by induction that the gauge is accurate in general, that is, that the gauge is reliable.<sup>3</sup>

The important thing to bear in mind is that every step Roxanne takes is sanctioned by reliabilism. Using ' $K$ ' to stand for 'Roxanne knows that...', we can represent her procedure as:

- (1)  $K$ (Tank is full at  $t_1$ ) Reliable Process
- (2)  $K$ (Gauge reads "F" at  $t_1$ ) Perception
- (3)  $K$ (Gauge reads "F" at  $t_1$  & Tank is full at  $t_1$ ) Logical Inference
- (4)  $K$ (Gauge reads accurately at  $t_1$ ) Logical Inference
- (5) Repeat<sup>4</sup>
- (6)  $K$ (Gauge is reliable) Induction

Of course, something is amiss. Roxanne cannot come to know that the gauge is reliable in this way. In trying to understand what is wrong, it is helpful to compare what Roxanne does with a similar but legitimate procedure. Say that Roxanne's friend Omar is interested in the reliability of the gas gauge in his car. Omar has a dipstick which he uses to measure the level of fuel in the tank. He compiles a "track-record" for the gauge just like Roxanne's.<sup>5</sup> I assume that what Omar does is quite all right. He can establish that the gauge is reliable in the manner described. So, from the original result that Roxanne *does not* know that the gauge is reliable, that is,  $\neg(6)$ , it follows under reasonable assumptions that  $\neg(4)$ , and  $\neg(3)$ . Seemingly, (3) follows

<sup>3</sup>Note that Roxanne could continue and infer that, since her beliefs about the amount of fuel in the tank are due to a reliable gauge, her beliefs themselves are formed in a reliable way. See below.

<sup>4</sup>Maybe a further step is needed to collect these results, so that there is an additional premise:

$K$ (Gauge reads accurately at  $t_1$  & Gauge reads accurately at  $t_2$  & ...).

I do not think that anything turns on this, so from now on, I will set it aside.

<sup>5</sup>For an account of "track-record arguments," see William Alston, *The Reliability of Sense Perception* (Ithaca: Cornell, 1996), p. 9.

impeccably from (2) and (1). Since (2) is not open to dispute, we arrive at  $\neg(1)$ . I will call this undoing of Roxanne's knowledge "rollback." The reliabilist has no resources to resist rollback. So if she concedes, as she should, that Roxanne's endeavors do not produce knowledge, contradiction threatens. The reliabilist's core doctrine commits her to (1), yet rollback would lead to the rejection of (1).

Roxanne does not know that her car's gas gauge is reliable. Why not? As stated, she has no information at all—no evidence—pertaining to whether the gauge is reliable or not. According to internalism, knowledge always requires justification, and justification requires evidence (at least in cases like this one).<sup>6</sup> Since Roxanne lacks evidence which would support the claim that the gauge is reliable, internalism has the consequence that Roxanne does not know that the gauge is reliable. The internalist gets the Gas Gauge Case right. By contrast, the reliabilist has to say that Roxanne knows that the gauge is reliable. The reliabilist gets the Gas Gauge Case wrong. No matter what else may be true, this difference is a strong consideration in favor of internalism and against reliabilism.

It would be a mistake to think that the Gas Gauge Case is a quirk, a superficial failing on the reliabilist's part to get all of the phenomenology exactly right. Rather, this case highlights and crystallizes what I take to be one of the fundamental difficulties reliabilism faces.<sup>7</sup> Consider the view I will call *simple reliabilism*:

(SR) *S* knows that *p* just in case, in the neighborhood of nearby possible worlds *N*, *S* does not believe falsely that *p*.

SR is clearly too weak. Whenever someone believes a metaphysically necessary or nomically robust proposition, that person will count as knowing according to SR. For example:

*Lucky Scientist Case.* Mitch is trying to determine the value for an important, fine-tuned physical constant. Unfortunately, his data are defective, but by luck he happens on the right answer, *C*. If the value of the constant were much different from *C*, the universe would be utterly unlike the way it actually is; say electrons would not bind to nuclei. So, there are

<sup>6</sup>This is just a stipulation of how I will use the term 'internalism'.

<sup>7</sup>And, I suspect, not just reliabilists. A virtue epistemologist will have to say that Roxanne's undertaking involves some lapse in epistemic virtue. How so? She uses perception and reasoning in appropriate ways. What she does not do is to restrict her belief to propositions for which she has the evidence she needs. That is a shortcoming on her part. But if that counts as a failure to manifest epistemic virtue, virtue epistemology seems to be collapsing into internalism. Cf. Noah Lemos, "Epistemic Circularity Again," in Ernest Sosa and Enrique Villanueva, eds., *Philosophical Issues*, Volume 14 (Malden, MA: Blackwell, 2004), pp. 254–70.

no nearby possible worlds in which the value of the constant is not *C*. It follows that there are no nearby possible worlds in which Mitch believes falsely that the value of the constant is *C*.

According to *SR*, Mitch knows that the value of the constant is *C*.

Sophisticated reliabilists recognize that an amendment is needed. They observe that drawing conclusions on the basis of defective data is an unreliable belief forming process; it produces many false beliefs in nearby worlds. *SR*, then, should be supplanted by some view along the lines of “process reliabilism”:

(*PR*) *S* knows that *P* just in case *S*'s belief that *P* is formed by a reliable process, where a reliable process is one which yields mostly true beliefs in the neighborhood *N* of nearby possible worlds.

*PR* properly handles the Lucky Scientist Case. Mitch's belief that the value of the constant is *C* does not amount to knowledge because it is the result of an unreliable belief forming process. So far, so good. But consider some reliable belief forming process, for example, using the gauge in Roxanne's car to determine the amount of fuel in the car's gas tank. Call this process (*G*).<sup>8</sup> Suppose further that someone, say Roxanne, believes that *G* is reliable. That is, she believes that *G* generates few false beliefs in *N*. Roxanne's belief satisfies the requirements of *SR*. The reliability of *G*, its tendency to produce true rather than false beliefs, is ultimately a nomic fact. In all the possible worlds in *N*, if someone believes that *G* produces few false beliefs in *N*, what he believes will be true. In this way, Roxanne's belief that *G* is reliable is like Mitch's belief that the value of the physical constant he is investigating is *C*.

The proponent of *PR* will be quick to point out that Roxanne does not know that *G* is a reliable process unless she forms a belief to that effect *by a reliable process*. But if *G* is a reliable process, then beliefs about *G* formed by using *G* itself are beliefs formed by a reliable process. So, if Roxanne uses *G* (supplemented, as necessary, by some other reliable processes) to form the belief that *G* is a reliable process, the requirements of *PR* will be met. Roxanne will know that *G* is a

<sup>8</sup> I imagine Roxanne as proceeding in two stages. First, she arrives at a belief that the gas gauge is reliable, as above, and proceeding from there she concludes that the process by which forms beliefs about the fuel in the tank is reliable. Since she has no justification for the former, she has no justification for the latter. Roxanne could also omit the first step, and treat herself like the gas gauge, as it were. That is, she could combine her reliably formed belief that the gas tank is full, with her reliably formed belief that she believes (in some specified way) that the gas tank is full, and so on. She would then infer that her way of forming beliefs about how much fuel is in the tank is reliable. See below.

reliable process. And, of course, that is just how the Gas Gauge Case goes. Roxanne arrives at her belief that *G* is reliable using only reliable processes, but in doing so she does not acquire any evidence that *G* is reliable. The reliabilist's criteria for knowledge are satisfied, while the internalist's are not. This is no surprise. Arriving at a belief by way of a reliable process is one thing, and acquiring evidence is something else again. Given this difference, there is nothing in the nature of internalism to suggest that it will go wrong in the same situations, for the same reasons, that reliabilism goes wrong.

## II. TU QUOQUE?

My view is that the Gas Gauge Case exposes a serious defect which burdens reliabilism, but not internalism. In principle, there are three lines of response. First, what appears to be a defect is really no such thing. Bootstrapping is acceptable, the Gas Gauge Case notwithstanding.<sup>9</sup> For some consideration of a view along these lines, see the Appendix. Second, reliabilism does not have the defect I claim it does. Properly understood or appropriately revised, reliabilism does not sanction bootstrapping after all.<sup>10</sup> I will not take up this possibility here. Finally, as Stewart Cohen and James Van Cleve have argued, bootstrapping is objectionable, but there is a broad difficulty involving bootstrapping which confronts internalism as well as reliabilism.<sup>11</sup> I will examine this claim in detail. But it is important to see at the outset that, even if Cohen and Van Cleve are right, the relative shortcomings of reliabilism brought to light by the Gas Gauge Case will not be eliminated.<sup>12</sup> However other things may turn out, reliabilism gets certain things wrong that internalism gets right. Thus, internalism is better off than reliabilism.

<sup>9</sup> See Alston (*op. cit.*), with reservations; James Van Cleve, "Is Knowledge Easy or Impossible? Externalism as the Only Answer to Skepticism," in Steven Luper, ed., *The Skeptics* (Aldershot, UK: Ashgate, 2003), pp. 45–59, with reservations; Michael Bergmann, "Epistemic Circularity: Malignant and Benign," *Philosophy and Phenomenological Research*, LXIX (2004): 709–27; and Frederick Schmitt, "What Is Wrong with Epistemic Circularity?" *Philosophical Issues*, Volume 14 (*op. cit.*).

<sup>10</sup> See Jose Zalabardo, "Externalism, Skepticism, and the Problem of Easy Knowledge," *Philosophical Review*, cxiv (2005): 33–61. For objections to Zalabardo's view, see Anthony Brueckner and Christopher Buford, "Bootstrapping and Knowledge of Reliability" (in preparation).

<sup>11</sup> Stewart Cohen, "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research*, LXV (2002): 309–29, and Van Cleve, *op. cit.*

<sup>12</sup> Here Cohen and Van Cleve may differ somewhat. Cohen holds that reliabilism and internalist foundationalism share a common a common commitment and a common defect, namely allowing bootstrapping. To that extent, the fact that reliabilism sanctions bootstrapping is the basis of an objection to the view, all the same. Van Cleve appears more inclined to think that reliabilism's permitting bootstrapping ultimately does not count against it.

Let me begin with Cohen's account, which will take some unpacking. He tells us that there is a problem of "easy knowledge" which arises for both reliabilism and a certain version of internalism: "The problem is that ... we can acquire reliability knowledge very easily—in fact, all too easily from an intuitive perspective ... We can call this the 'Problem of Easy Knowledge'" (*op. cit.*, p. 311). Cohen provides different accounts of what "reliability knowledge" is (*op. cit.*, p. 311). He writes at one point: "There is nothing to stop us from acquiring by trivial inferences, all sorts of knowledge about how we are not deceived or misled by our belief sources" (*op. cit.*, p. 315). This is knowledge that various *error-possibilities* do not obtain. These include that the possibility that the table that looks red to you is in fact a white table illuminated by red lights and the full-blown skeptical possibility that you are a thoroughly deceived brain in a vat. Cohen alludes to another notion of reliability, and concomitantly another sort of reliability knowledge. A belief-forming process like color-vision is reliable in this sense (*op. cit.*, pp. 316–17). *SR* and *PR* discussed above are attempts to analyze what such reliability amounts to. When I use the term 'reliability,' I will be talking about this second conception.

These points indicate that there are really two distinct sets of issues that Cohen has brought together under the rubric of "easy knowledge," and he would agree.<sup>13</sup> One of these explicitly has to do with bootstrapping. According to Cohen, bootstrapping generates "easy knowledge." The other set of issues has to do with the application of the Closure Principle for Knowledge.<sup>14</sup> The Moorean response to skepticism is an illustration. Moore said that he knew by seeing it that he had a hand in front of him. Given the Closure Principle for Knowledge, Moore thereby knew that he was not deceived by an evil demon. Cohen calls this, too, "easy knowledge." We now need to guard against a certain confusion. Suppose that internalism and reliabilism both run into trouble with the Closure Principle for Knowledge, as Cohen maintains (*op. cit.*, pp. 312–16). In that sense, both views have a problem about "easy knowledge." But it would be a mistake to go on to say that, since the difficulty for reliabilism illustrated by the Gas Gauge Case is also subsumed under the rubric of "easy knowledge," and since both reliabilism and internalism have *some* difficulty pertaining to "easy knowledge," the two views are equally good or bad. Whether

<sup>13</sup> Cohen does believe that the two problems have a common root, or, perhaps it is better to say, a common solution. The solution would be to reject the principle *KR*; see below.

<sup>14</sup> The Closure Principle for Knowledge says, roughly, that if *S* knows that *P*, and also knows that *P* entails *Q*, then *S* knows, or can know, *Q*.

reliabilism and internalism have trouble with the Closure Principle for Knowledge is irrelevant to how they stand with respect to bootstrapping.

Cohen and Van Cleve do believe that the internalist is committed to at least some bootstrapping. Van Cleve writes:

But evidentialism as such is not enough to thwart bootstrapping. An evidentialist view will also permit bootstrapping, just so long as it says there are sources that confer knowledge or justification without any requirement that the reliability of the source be known .... In fact, in order for the bootstrapper to get off the ground, all we need allow is that perception is a source of knowledge independent of knowledge of its own reliability, and similarly for introspection, memory, and induction. For by introspection one may know that perception is the source of one's belief in *p*; by perception ... itself one may know *p*; by combining these two, one may know that perception was truthful in a given instance; by memory one may then know that perception was truthful in a great many instances (and deceptive, we may assume, in few or no instances); and by induction, one may then conclude that perception is reliable. It seems that one could take issue with this result only by questioning the sufficiency of perception, introspection, memory or induction to give knowledge—or at least their sufficiency in the absence of any knowledge of their reliability (*op. cit.*, pp. 48–49).

Cohen makes the same point by way of an example:

According to that view [a version of internalism—JV], I can know the table is red on the basis of its looking red, even though I have no prior evidence that something's looking red is a reliable indication that it is red. But then once I know the table is red, I can appeal to that fact in reasoning. A little introspection will tell me that the table appears red. So now I know that the table looks red and that it is red. So I now have some evidence that something's looking red is a reliable indication that it is red. And by taking a few more looks, I can acquire more evidence (*op. cit.*, p. 317).

The target here is internalism, or more particularly a foundationalist strain of internalism. According to this view, there are basic beliefs generated by memory, introspection, and perhaps perception. Experiences of the appropriate type serve as evidence which immediately justifies these beliefs, but the beliefs are still basic, in that they are not justified by any other *beliefs*.<sup>15</sup> That is, no other belief is epistemically prior to memory, introspective, or perceptual beliefs.

<sup>15</sup> This is one prominent conception, among others, of how basic beliefs are justified. For discussion, see Vogel, "Internalist Responses to Skepticism," in John Greco, ed., *The Oxford Handbook of Skepticism* (New York: Oxford, forthcoming). My sense is that the treatment in the text applies to the other conceptions, so long as they provide that there is some rule or other that specifies which beliefs are basic.

I have characterized the view of interest somewhat more narrowly than Cohen and Van Cleve do. They are particularly concerned with the following epistemic principle:

(*KR*) We can know a deliverance of a “belief source” *F* only if we first know that *F* is reliable.<sup>16</sup>

“First” in this context denotes epistemic rather than temporal priority. The reference to “belief sources” is problematic in a number of ways, although it is fairly standard.<sup>17</sup> Cohen and Van Cleve group together all views which reject *KR*. These include reliabilism and the sort of internalist foundationalism just described. Internalist foundationalism must deny *KR*, because a basic belief is by definition one that is not supported by any other epistemically prior belief. If *KR* were true, even a supposedly basic belief would require support from a reliability belief. The supposedly basic belief would not be basic after all. It is worth observing that *KR* admits of some variants:

(*KR*+) In *all* cases, we can know a deliverance of a “belief source” *F* only if we first know that *F* is reliable.

(*KR*<sup>o</sup>) In *some* cases, we can know a deliverance of a “belief source” *F* only if we first know that *F* is reliable.

(*KR*-) In *no* case, does our knowing a deliverance of a “belief source” *F* require us to know first that *F* is reliable.

The differences among these principles shed some light on why the Gas Gauge Case is a problem for the reliabilist, but not for internalists. *KR*+ is just a slightly more explicit version of *KR*, and we have seen why an internalist foundationalist has to reject it. *KR*- captures the reliabilist’s position. But all internalists, including the internalist foundationalist, may accept *KR*<sup>o</sup>.<sup>18</sup> Internalists are thus in a position to

<sup>16</sup> Cohen originally presents the weaker principle, “A potential knowledge source *G* can yield knowledge for *S*, only if *S* knows *G* is reliable,” which does not include any reference to epistemic priority (*op. cit.*, p. 308).

<sup>17</sup> Note that *KR* is supposed to block bootstrapping, but in the Gas Gauge Case, Roxanne’s bootstrapping had to do with the reliability of her gas gauge. The gauge is a “belief source” only in a very stretched sense. Another concern is that loose talk about “belief sources” may bring on the generality problem for reliabilism. Suppose I am in the supermarket ready to buy some groceries and decide to carry out a little bootstrapping. I remember or take myself to remember the items on the shopping list. What is my “belief source”? Is it my memory for items on shopping lists, my memory of lists in general, or general memory as such? Finally, Cohen and Van Cleve mean induction to fall under the scope of *KR*, but induction hardly seems like a belief source in the way that perception, memory, or introspection are. This discrepancy suggests that something significant has escaped from view. See discussion in section III, below.

<sup>18</sup> The internalist foundationalist can endorse *KR*<sup>o</sup> so long as the belief source involved is not one that is supposed to generate basic beliefs.



say, appropriately, that Roxanne does not know that her car's gas tank is full, because she does not know that her "belief source" (that is, going by the gas gauge) is reliable. Roxanne lacks this knowledge by internalist lights, because she has no pertinent evidence about the gauge. Internalism is thus more discriminating than reliabilism. The internalist can, in principle, distinguish cases where knowledge of the reliability of one's belief source is required for knowledge from cases where such knowledge is not required. The reliabilist's commitment to *KR*— prevents him from doing the same. In this important way, internalism is superior to reliabilism.

What, more generally, does *KR* have to do with bootstrapping? Let ' $X_i$ ' stand for the various propositions known by way of belief source  $F$ , let ' $B_F$ ' stand for ' $F$  yields the belief that ...', and let ' $RF$ ' stand for ' $F$  is a reliable belief source'. Bootstrapping generates a track-record argument:<sup>19</sup>

- (1)  $X_i$
- (2)  $B_F(X_i)$
- (3)  $B_F(X_i) \ \& \ X_i$
- (4)  $B_F$  yields a true belief with respect to  $X_i$
- (5) Repeat
- (6)  $RF$

Obviously, if bootstrapping generates justification for  $RF$  in this way, then  $X_i$  is epistemically prior to  $RF$ . But, according to *KR*,  $RF$  is epistemically prior to  $X_i$ . The result is that  $X_i$  is prior to  $RF$  and  $RF$  is prior to  $X_i$ . That is impossible, because epistemic priority is anti-symmetric. Hence, *KR* is incompatible with bootstrapping, and views that embrace *KR* will thereby avoid bootstrapping.

This outcome is hardly encouraging, though, because *KR* apparently leads to skepticism. Van Cleve puts the argument succinctly:

- (1) (*KR*) We can know that a deliverance of a belief source  $F$  is true only if we first know that  $F$  is reliable.
- (2) We can know that  $F$  is reliable only if we first know, concerning certain of its deliverances, that they are true.<sup>20</sup>

<sup>19</sup> Undoubtedly, we often arrive at reliability judgments in much more complicated and subtle ways, but I will join others in assuming that track-record arguments are all right so far as they go. In a way, though, talking about arguments at all is somewhat misleading. The real concern is the structure of one's knowledge or justified beliefs, which is not the same thing. But, with that caveat, I will continue to examine the structure of various arguments, taking for granted that the points made will carry over to analogous results about knowledge or justified belief.

<sup>20</sup> Following Van Cleve, *op. cit.*, p. 50.

*F* provides knowledge only if it satisfies the conditions given by (1) and (2). Let us grant (2). Combined with (1), it gives us the result that some deliverances of *F* are epistemically prior to themselves, which is impossible. Hence, *F* does not satisfy (1) and (2), and beliefs due to *F* are not knowledge. The argument is perfectly general, and applies to any belief source *F*. The upshot is that we have no knowledge at all.<sup>21</sup>

Cohen and Van Cleve go on to consider whether some combination of *KR* with internalism, or with an internalist-reliabilist hybrid, might keep both bootstrapping and skepticism at bay, but they find that the outlook is bleak. In any event, their view is that internalist foundationalism, like reliabilism, is inevitably tainted with bootstrapping of some sort.

### III. CUTTING THE BOOTSTRAPS

So far, the discussion of bootstrapping has been conducted in terms of knowledge, but insofar as the prospects for internalism are our present concern, it will be more straightforward to talk about justification instead. There is an analogue of *KR* for justification:

(*JR*) We are justified in believing a deliverance of a belief source *F* only if we first have a justified belief that *F* is reliable.

And, corresponding to the Closure Principle for Knowledge, there is a Closure Principle for Justification.<sup>22</sup>

A further preliminary point needs to be made. As I indicated earlier, we should be wary of framing the issues connected with bootstrapping in terms of the outputs of “belief sources.” A preferable approach is to distinguish between *indirect* and *direct* justification. Consider David, who sees some fresh rabbit tracks in the snow, and forms the belief that (RABBIT) a rabbit has recently been at that spot. His evidence is that (TRACKS) there are tracks in the snow of a distinctive shape. David could not have a justified belief that RABBIT on the basis of TRACKS alone. He also needs to have a justified belief that tracks with the distinctive shape he sees are present only if they are made recently by rabbits. I will say that, for David, the evidential relation between TRACKS and RABBIT is indirect. Contrast this situation with the way the foundationalist conceives of basic beliefs. Say you have a basic introspective belief that (SLEEPY) you feel sleepy. This belief is justified by (EXP) your experience of feeling sleepy. EXP is sufficient

<sup>21</sup> I reject both premises and the conclusion. For discussion, see Vogel (forthcoming).

<sup>22</sup> Roughly, if *S* has a justified belief that *P*, and knows that *P* entails *Q*, then *S* is justified in believing *Q*.

by itself to justify SLEEPY. You do not need to have any additional epistemically prior justified belief that if EXP occurs, then SLEEPY. In this case, I will say that the evidential relation between EXP and SLEEPY is direct.<sup>23</sup>

The distinction between direct and indirect evidential relations helps us sharpen the challenge to internalism posed by Cohen and Van Cleve. Take a stock example of induction, for example one in which Ira infers that (AC) all samples of copper conduct electricity from (OC) repeated observations of copper samples. OC is Ira's evidence for AC, and Ira's justification for believing AC is direct. Ira does not need an ancillary justified belief that OC implies AC in order to be justified in believing AC on the basis of OC. Now we can say what basic beliefs gained by introspection, memory, or perception have in common with beliefs due to induction. These beliefs are all directly justified by the evidence that supports them. I take Cohen and Van Cleve's view to be that, in these circumstances, the internalist has no resources to preclude bootstrapping.<sup>24</sup>

Cohen and Van Cleve mean to confront the internalist with a dilemma: Either accept JR or countenance bootstrapping. However, neither author presents an argument that the alternatives on offer are exhaustive, so there is a lacuna in their position. It is not obvious at the outset that accepting JR is the only route by which the internalist could avoid bootstrapping. Hence, it seems premature to claim that internalism shares any defect with reliabilism.

It would be nice to be able to settle the issue by showing how internalism can avoid bootstrapping without accepting JR. In this vein, we might begin by asking why the internalist, noting the impropriety of bootstrapping, cannot reasonably assert that there is an epistemic principle which forbids it, namely: No bootstrapping! The trouble is that "bootstrapping" is a term of art, introduced by examples. So,

<sup>23</sup> To be more explicit: If *S* has evidence for *E* for *H*, but *S* would not be justified in believing *H* on the basis of *E* without also having a prior justified belief for which she needs additional evidence, then the evidential relation between *E* and *H* is *indirect*. Otherwise, the evidential relation between *E* and *H* is *direct*. This distinction is close to, but not quite the same as, the distinction between mediately and immediately justified belief.

<sup>24</sup> Note that foundationalism in the traditional sense is not really at issue here. Even if the internalist were to abandon the claim that there are basic beliefs, bootstrapping would still be a problem with respect to induction, at least. See below. To be fair to Cohen and Van Cleve, they will say that inductive beliefs along with introspective, memorial, and perceptual beliefs count as "basic" in their special sense, insofar as those beliefs can count as knowledge without satisfying KR. In addition, Cohen raises the concern that coherentist views may foster their own special version of bootstrapping (*op. cit.*, p. 323); see also Schmitt (*op. cit.*, pp. 388–89). If so, that is all the more reason why internalists of all stripes need some account of why bootstrapping is not permitted.

without further clarification, we cannot do much to evaluate the suggested prohibition or its consequences.

To make progress, it is helpful to bear in mind that Roxanne's procedure in the Gas Gauge Case is, as Van Cleve puts it, a parody of the *Inductive Justification of Induction* (hereafter '*IJI*').<sup>25</sup> A standard formulation of *IJI* goes as follows:

- (1) Induction yields the conclusion  $X_i$ .
- (2)  $X_i$ .
- (3) Induction yields the conclusion  $X_i$  and  $X_i$ .
- (4) Induction yields a true belief with respect to  $X_i$ .
- (5) Repeat.
- (6) Induction is reliable.

*IJI* is typically formulated in terms of next case induction.<sup>26</sup> Inductive inferences of this sort would be:

Every observed lemon is yellow, so the next lemon I encounter will be yellow

and

Every observed sapphire is blue, so the next sapphire I encounter will be blue.

Moreover, as *IJI* is standardly presented, these inductive conclusions  $X_i$  are *verified in some other way*. For example, I find another lemon, and check by perception that it is yellow, or I find another sapphire and check by perception that it is blue.

Now, suppose that I do not bother to perform these checks. We then have what deserves to be called the *Lazy Justification of Induction* (hereafter '*LJI*'):

- (1) Every observed lemon is yellow.
- (2) Hence, all lemons are yellow.
- (3) The belief that all lemons are yellow was generated by induction.
- (4) Induction yields the belief that all lemons are yellow, and all lemons are yellow (from 2, 3).
- (5) That is, induction yields a true belief in this instance.
- (6) Every observed sapphire is blue.
- (7) Hence, all sapphires are blue.
- (8) The belief that all sapphires are blue was generated by induction.

<sup>25</sup> Van Cleve, *op. cit.*, p. 47.

<sup>26</sup> See, for example, Max Black, "Self-Supporting Inductive Arguments," reprinted in Richard Swinburne, ed., *The Justification of Induction* (New York: Oxford, 1974), pp. 127–34.

- (9) Induction yields the belief that all sapphires are blue, and all sapphires are blue (from 7, 8).
- (10) That is, induction yields a true belief in another instance.
- (11) Repeat.
- (12) Induction is reliable.

Not even proponents of the standard inductive justification of induction will happily accept this line of thinking.

The bootstrapping Cohen and Van Cleve try to foist on the internalist follows the same pattern as *LJI*. Suppose you have a directly justified belief  $H_i$  for which the evidence is  $E_i$ . Then there is a legitimate epistemic principle, a rule, which says something like ‘if you have evidence  $E_i$ , then believe  $H_i$ ’, or perhaps ‘if you have evidence  $E_i$ , then you are justified in believing  $H_i$ ’. Here is the *Rule Bootstrapping Argument (RBA)*:

- (1)  $E_i$
- (2)  $H_i$  (using the rule  $R$ ).
- (3) Rule  $R$  generates the conclusion  $H_i$ .

Note that, just as with *LJI*, there is no reason to believe  $H_i$  beyond its support from  $E_i$  via rule  $R$ .

- (4) Rule  $R$  generates the conclusion  $H_i$  &  $H_i$  (from 2, 3).<sup>27</sup>
- (5) Repeat.
- (6) Therefore, the use of rule  $R$  is reliable.

Now, the status of *IJI* has been controversial. Its detractors say that the argument is viciously circular. In response, some of its defenders have distinguished

*Premise Circularity*: The conclusion of the argument is a premise of the argument

from

*Rule Circularity*: An epistemic rule  $R$  is employed in an argument for the conclusion that the use of rule  $R$  is reliable.<sup>28</sup>

<sup>27</sup> At this point, one might add the premise that rule  $R$  generates a true belief with respect to  $H_i$ , as I have done above. I discuss this point in section iv.

<sup>28</sup> There are different formulations of what rule-circularity is, although the one I have given here is fairly common. Cf. Van Cleve, “Reliability, Justification, and the Problem of Induction,” in Peter French, Theodore Uehling, and Howard Wettstein, eds., *Midwest Studies in Philosophy*, Volume 9 (*Causation and Causal Theories*) (Minneapolis: Minnesota, 1984), pp. 555–67, see p. 558. At any rate, it has its shortcomings. For example, suppose that one used rule  $R$  to argue for the conclusion that either  $R$  is reliable or it will be sunny tomorrow. Strictly speaking, the path to that conclusion is not rule-circular, although it does seem defective. My sense is that a good deal more needs to be done to improve our understanding of rule-circularity.

Everyone agrees that premise-circularity is unacceptable, but *III* is not premise-circular. The conclusion of *III*, that induction is reliable, does not appear as a premise of the argument. *III* is rule-circular, however. The rule of inductive inference which permits generalization from observed cases to unobserved cases is used to reach the conclusion that the application of that rule (that is, induction itself) is reliable. Proponents of *III* hold that rule-circularity is benign. Consideration of *LII* strongly suggests otherwise. The claim that induction is reliable is never used as a premise in *LII*, so *LII* is not premise-circular.<sup>29</sup> But *LII* is clearly rule-circular.<sup>30</sup> Premise (2) is licensed by the application of *R*, and (2) is then enlisted to support the conclusion that the use of *R* is reliable. Rule-circularity is unacceptable, at least in contexts like this one.

These observations suggest that the internalist can prevent bootstrapping by forbidding rule-circular justification:

(*NRC*) A belief that an epistemic rule *R* is reliable cannot be justified by the application of *R*. That is, neither the conclusion itself nor any belief which supports the conclusion may be justified in virtue of the application of *R*.

*NRC* is in somewhat the same vein as a proposal due to Richard Fumerton.<sup>31</sup> One way of reading Fumerton's remarks is that there is a prohibition against "self-support," giving us a "no self-support" principle:

(*NSS*) One cannot obtain ... a justified or warranted belief that a belief source *S* is trustworthy by relying even in part on source *S*.<sup>32</sup>

Setting aside the differences between the application of epistemic rules, on the one hand, and "belief sources" on the other, *NSS* seems to be strictly stronger than *NRC*. Suppose that you see something odd, like a

<sup>29</sup> One way to capture what separates Cohen and Van Cleve from me on these issues is that they take the defect in bootstrapping to be premise-circularity (if anything), while I take it to be rule-circularity. Many authors, following Alston, "Epistemic Circularity," reprinted in Alston, *Epistemic Justification: Essays in the Theory of Knowledge* (Ithaca: Cornell, 1989), pp. 319–49, are concerned with the acceptability of "epistemic circularity," which is variously understood.

<sup>30</sup> Rule-circularity obtains when the rule *R* is used to generate a premise of the argument or when it is used to move from premises of the argument to a conclusion. Discussions of *III* have focused on the second aspect, while my concern is with the first.

<sup>31</sup> He writes at one point: "You cannot *use* perception to justify the reliability of perception! You cannot *use* memory to justify the reliability of memory! You cannot *use* induction to justify the reliability of induction! Such attempts to respond to the skeptic's concerns involve blatant, indeed pathetic, circularity" (Fumerton, *Metaepistemology and Skepticism*, p. 177). I do not know whether Fumerton would agree that rule-circularity in particular is the problem. In "Epistemic Internalism," Fumerton has tempered his position, although he stands by the sentiments just quoted so far as they go.

<sup>32</sup> The idea is due to Michael Bergmann, "Externalism and Skepticism," *Philosophical Review*, cix (2000): 159–94, see p. 168, although the wording is Cohen's (*op. cit.*, p. 319).

gorilla waiting at a bus stop. You might look again, to make sure you saw what you think you saw. In this case, the agreement between the two deliverances of your perceptual belief source does seem to count in favor of the “trustworthiness” of the belief source that produced them. I suppose that such an appeal to coherence does rely “in part” on the source of the beliefs, and so runs afoul of *NSS*. But one can be justified in holding that two beliefs are both true because they agree, and therefore justified in holding that their source *S* is “trustworthy,” without the application of any rule which licenses beliefs due to *S*. Hence, justification by coherence is allowed under *NRC*, but apparently not by *NSS*.

To sum up, Cohen and Van Cleve maintain that internalists cannot avoid bootstrapping. But *NRC* seems like a tenable principle governing epistemic justification. If it is, then bootstrapping is impermissible by the internalist’s lights, and the Cohen-Van Cleve *tu quoque* does not come off.

#### IV. THE ROLLBACK PROBLEM

I would be very pleased to leave matters at that, but more needs to be said. The attempt to foreclose bootstrapping by *NRC* faces a serious difficulty. Consider *RBA* again. According to *NRC*, (6) is unjustified. The trouble is that this result seems to initiate a “rollback” of justification (see section 1). Consider the step from (4) to (6).<sup>33</sup> We might well think that if (4) is justified, so is (6), and, conversely, if (6) is not justified, then neither is (4). Then, since *NRC* denies justification to (6), (4) must be unjustified. Things get worse. (2) and (3) entail (4). If (2) and (3) are justified, then presumably (4) is, too. But we just said that (4) is unjustified. Hence, either (2) or (3) is unjustified. There seems to be no basis whatsoever for denying justification to (3). Hence, (2) must be unjustified. This is a disastrous result. Using *NRC* to head off bootstrapping seems to lead to skepticism with respect to the direct justification of any belief by any epistemic rule.<sup>34</sup>

A comparison between a more concrete example of bootstrapping and a foil may be helpful here. Suppose you are at the grocery store and decide to do some bootstrapping about memory. You remember that ( $X_i$ ) you are out of various things.<sup>35</sup> By hypothesis, your recall

<sup>33</sup> Of course, the justification for (6) would have to come from a multitude of conjunctions on the pattern of (4). For our purposes, these are all alike; what holds for one, holds for the others. So, for ease of exposition, I will write from now on as though (4) alone were all that *RBA* requires.

<sup>34</sup> *LJI*, in particular fits the pattern of *RBA*. So, by the present reasoning, the failure of *LJI* as a defense of induction is enough to establish inductive skepticism.

<sup>35</sup> Ordinarily ‘remember’ is factive. What I mean here is that you seem to remember that you need to buy certain things. When I use the terms ‘remember’ and ‘memory’, I will be talking about apparent remembering and apparent memory.

gives you evidence  $E_i$  which supports your belief  $X_i$ . You are aware by introspection that you remember  $X_i$ , which I will abbreviate as ‘Mem ( $X_i$ )’. You then bootstrap to the conclusion that your memory is reliable:

- (M1)  $E_i$
- (M2)  $X_i$
- (M3) Mem( $X_i$ )
- (M4) Mem ( $X_i$ ) &  $X_i$
- (M5) Memory is reliable

Instead of bootstrapping, however, you could have checked your memory against a shopping list. Consulting the list gives you evidence that you are out of various items, just as memory-experience does. But consulting the list does more. In particular, referring to the list plausibly does give you evidence that your memory is reliable.<sup>36</sup> If this is correct, it sheds some light on the rollback problem. My suggestion is that your evidence  $E_i$  supports (M2) and (M4), but not (M5). Hence, you have justification for (M2) and (M4), but not (M5). Note that, so far, this assessment is consistent with the Closure Principle for Justification, in that (M4) does not entail (M5).

The trouble is that, on the account just given, (M4) is supposed to be justified, (M4) seems to justify (M5), yet (M5) is supposed to be unjustified. I can see two ways to address this issue. According the first approach, (M4) is insufficient to justify the reliability claim (M5). To be justified in believing that your memory is reliable, you also need justification for believing that your memory is not *mistaken*, that is,  $\neg(\text{Mem}(X_i) \ \& \ \neg X_i)$ .<sup>37</sup> If your memory did frequently get the facts wrong, it would hardly be reliable. The thought here is that a memory experience  $E_i$  may justify  $X_i$ , and also contribute to justifying the conjunction  $(\text{Mem}(X_i) \ \& \ X_i)$ , but  $E_i$  does not justify  $\neg(\text{Mem}(X_i) \ \& \ \neg X_i)$ .<sup>38</sup> To establish that your memory is not mistaken requires something more, like checking what you remember against the shopping list. When the two match, you acquire evidence that  $\neg(\text{Mem}(X_i) \ \& \ \neg X_i)$ , which in turn provides essential support for the claim that your memory is reliable. Without an additional contribution of this sort, (M4) does

<sup>36</sup> Not just memory, in fact, but that plus introspection or whatever discloses that you remember that you need the various items.

<sup>37</sup> Van Cleve says “we can assume” that the evidence available to the bootstrapper includes evidence that memory has “been deceptive on few occasions” (“Is Knowledge Easy or Impossible?” p. 49). The approach under discussion questions that assumption.

<sup>38</sup> Why not? Some may find the claim intuitive as it stands. Also, since  $(\text{Mem}(X_i) \ \& \ \neg X_i)$  entails  $\text{Mem}(X_i)$ ,  $\text{Mem}(X_i)$  is positively relevant in the probabilistic sense to  $(\text{Mem}(X_i) \ \& \ \neg X_i)$ . On certain views about confirmation theory, it follows that  $\text{Mem}(X_i)$  is not evidence for  $\neg(\text{Mem}(X_i) \ \& \ \neg X_i)$ .



not justify (M5).<sup>39</sup> If (M4) as such does not justify (M5), the internalist can deny justification to (M5), without initiating the rollback of justification for other beliefs.

There are several points to note about this proposal. First, it stands on its own as an account of why and how bootstrapping fails, independently of *NRC*. Second, it may intensify the problem of induction. We might hope that we are justified, somehow, in believing that induction is reliable. On the present line of thought, that would require us to have evidence that our observations are not pointing us towards wrong conclusions. It is hard to know what such evidence would be, or how one could possibly acquire it.<sup>40</sup> Finally, this way out of the rollback problem has a significant cost, at least to some. The idea is that you lack justification for (M5), because you lack justification for  $\neg(\text{Mem}(X_i) \ \& \ X_i)$ . Now,  $X_i$ , that is, (M2), entails  $\neg(\text{Mem}(X_i) \ \& \ \neg X_i)$ . By the Closure Principle for Justification, if you are justified in believing (M2), you are justified in believing  $\neg(\text{Mem}(X_i) \ \& \ \neg X_i)$ . Conversely, if you lack justification for  $\neg(\text{Mem}(X_i) \ \& \ \neg X_i)$ , as proposed, then you must lack justification for (M2). But that is the unacceptable result we were trying to ward off in the first place. We see that the current proposal can succeed only if the Closure Principle for Justification fails.

An alternative approach to the rollback problem is to say that the rule-circularity of the support for (M5) vitiates the justification that would otherwise accrue to (M5) from (M4). Other things being equal, (M4) justifies (M5). What we have then is just an ordinary, acceptable track-record justification for (M5). However, (M4) together with

(RC) The support for (M5) is rule-circular

does not justify (M5). *RC* is a *defeater* of the justification (M4) provides to (M5), in just the sense that (M4) and *RC* together *do not* justify (M5), even though (M4) alone would justify (M5).<sup>41</sup> Because the

<sup>39</sup> Van Cleve writes: "One cannot help but think it relevant that in the case of the gauge, there is the possibility of making independent checks .... In the case of an ultimate source of knowledge such as perception or memory, however, there is no such possibility of an independent check (and no hope for support except self-support)" ("Is Knowledge Easy or Impossible?" p. 57). My explanation of why corroboration like this is needed must differ from Van Cleve's, though. See note 37.

<sup>40</sup> The new riddle of induction is uncomfortably nearby.

<sup>41</sup> Some philosophers, notably John Pollock, distinguish a special class of undercutting defeaters. On one way of understanding this notion, "Undercutting defeaters attack the connection between the reason and the conclusion rather than attacking the conclusion directly." See Pollock, "OSCAR: A General Theory of Rationality," reprinted in Robert Cummins and Pollock, eds., *Philosophy and AI: Essays at the Interface*, (Cambridge: MIT, 1991), pp. 209–26, on p. 209. *RC* would be a defeater of this sort. *RC* undoes the efficacy of (M4) as a reason to believe (M5), but *RC* is not a reason to deny (M5).

justification (M4) provides for (M5) is defeated, your lack of justification for (M5) does not imply a lack of justification for (M4).<sup>42</sup> Roll-back is blocked, the threat of skepticism is dissipated, and the Closure Principle for Justification is unviolated.<sup>43</sup>

Apart from its having these attractive consequences, what can be said in favor of this second proposal? An example suggests that *RC* can work as a defeater for a reliability claim:

*Record Keeping Case.* While you are bootstrapping in the grocery store, you keep a record of the results. Your notes have entries like “I remember that we need catfood/Yes, we do need catfood.” You go over your notes sometime later. At that point, you mistakenly but justifiedly believe them to record an episode when you were not bootstrapping, but rather checking your memory by reference to the shopping list instead (say the notes were misfiled in the “nonbootstrapping” folder).

Under these circumstances, I think, you have (some) justification for believing that your memory is reliable. Now imagine that you find out later that your notes actually record the results of bootstrapping. You realize that you employed a rule like ‘If you seem to recall  $X_i$ , then (you are justified in believing)  $X_i$ ’ to generate the entries, which in turn serve to justify your belief that employment of the rule is reliable. That is, you discover that the support for your reliability belief is rule-circular.<sup>44</sup> I suppose that, in this situation, your justification for the reliability belief would erode. Your belief that *RC* would defeat your justification for the belief that the rule you used is reliable.

<sup>42</sup> Compare: ‘Tweety is a bird’ justifies ‘Tweety flies’, but not given the defeater ‘Tweety is a penguin’. Your lack of justification for ‘Tweety flies’, given the defeater, does not imply that you lack justification for ‘Tweety is a bird’.

<sup>43</sup> An approach similar in spirit to the present one would be to note that the relation of justification or epistemic support is not (always) transitive. Accordingly, one might claim that although your original evidence including  $E_i$  supports (M4), and (M4) supports (M5), your original evidence does not support (M5). This account has various advantages and disadvantages, but I cannot pursue the matter here.

<sup>44</sup> Unfortunately, the example is less than ideally straightforward. Before the discovery, your justification was not rule-circular. You then had a justified belief that your memory had been favorably compared with the shopping list, which was good reason to believe that your memory is reliable. Afterwards, you no longer have justification for believing that this favorable comparison took place, and to that extent you are no longer justified in believing that your memory is reliable. But why does the discovery make (your awareness of) the notes thoroughly incapable of justifying your reliability-belief? Answer: You realize that you would in some way be using deliverances of your memory (involving applications of the rule), to justify the belief that your memory (the rule) is reliable. I am not sure whether this is rule-circularity *au pied de la lettre*, but I think it is close enough to make the point. See note 28.

There is another way to analyze this example, though. On this alternative line of thought, (M4) supports (M5). You would be justified in believing that your memory is reliable if you really were justified in believing (M4). A genuine comparison between your memories and the shopping list would provide evidence for (M4). The discovery that you really did not undertake such a comparison removes support from (M4), and thus removes support from (M5). This is the way, in general, that rule-circularity undoes the justification for reliability claims. If rule-circularity precludes justification for (M5), it must (somehow) undo the justification for (M4). But then, if (M4) is unjustified, we have the rollback problem the internalist needs to avoid.

I think that the status of the rollback problem comes down to this: Does the rule-circularity of your justification for (M5) mean that your justification for (M4) is defective, or give you reason to think that (M4) is not true? If yes, then (M4) is unjustified, and rollback ensues. The issues here are by no means clear-cut, but I do not see that this answer is the right one. In the first place, the specific flaw in the justification for (M5) does not attach to (M4). The justification for (M4) is not rule-circular, even though the justification for (M5) is. Second, would your discovery of bootstrapping in the Record Keeping Case indicate that your original beliefs about what grocery items you needed were false or unsupported? It seems not. Why should a subsequent, untoward inference proceeding from those original beliefs affect their epistemic standing? More broadly, if rule-circularity or anything else blocks the justification of a reliability belief downstream, that seems irrelevant to the justification of one's beliefs upstream. If this assessment is correct, then rule-circularity can explain the unacceptability of bootstrapping without bringing on the rollback problem.

#### V. CLOSING OBSERVATIONS

I have proposed that the use of bootstrapping to justify a reliability belief is rule-circular, and therefore illegitimate. My account located the defect in bootstrapping at the transition from (M4) to (M5) specifically. But one might well feel that trouble occurs earlier, and this impression will make the appeal to *NRC* seem artificial and misdirected. Bootstrapping proceeds from the claim that the use of a rule has been successful on particular occasions to the conclusion that use of the rule is reliable in general. But, when bootstrapping, are you justified in believing that the application of the relevant rule has *ever* been successful? Do you really have justification for believing that your memory has yielded a true belief in any instance? It may appear that you lack such justification unless you do something like checking

the shopping list, and this is the reason why bootstrapping cannot proceed. At issue here is the epistemic status of:

(M4T) My memory that  $X_i$  is true (or veridical).

(M4T) is logically equivalent to

(M4)  $\text{Mem}(X_i) \ \& \ X_i$ .

I have maintained that you are justified in believing (M4). Perhaps one could argue there is some epistemically significant difference between (M4) and (M4T), such that (M4) counts as justified while (M4T) does not.<sup>45</sup> And, possibly, such a discrepancy has something to do with the defectiveness of bootstrapping. But, in any case, it would still have to be shown why you cannot bootstrap from (M4) to (M5) directly, bypassing (M4T). The questions raised in section III and section IV would remain pressing, and the conclusions of those sections, such as they are, would remain applicable.

Epistemic bootstrapping raises any number of thorny issues. I have argued that, at a minimum, reliabilism exacerbates the problem bootstrapping poses, leaving the internalist significantly better off than the reliabilist. More than that, internalism, especially when fortified by *NRC*, may be able to avoid bootstrapping altogether. If not, then one would have to suspect that bootstrapping is really just the reflection of some other problem, which runs deeper still.

JONATHAN VOGEL

Amherst College

#### APPENDIX

Suppose that the only real alternatives we have are skepticism and accepting the legitimacy of bootstrapping. If we are forced to choose, the latter may very well seem like the better of the two. Michael Bergmann defends such a view. He observes that bootstrapping initially strikes us as distasteful, but he holds that some uses of it are quite all right. Bergmann writes: “Anyone who accepts the foundationalist thesis ... should think that epistemically circular track record arguments needn’t be a bad thing” (*op. cit.*, p. 713). The term “epistemically circular track record arguments” covers the kind of bootstrapping we have been discussing. In a nutshell, Bergmann’s view is:

What we have ... are two different contexts in which a belief in the trustworthiness of a belief source is formed. The first is a context in which the

<sup>45</sup> Defending such a view will not be easy, but for some thinking in roughly this direction, see Scott Sturgeon, “Truth in Epistemology,” *Philosophy and Phenomenological Research*, LI (1991): 99–108. It does seem quite intuitive that some kind of cognitive ascent is both essential to bootstrapping and problematic in some way.

subject has doubt or is uncertain about the source's trustworthiness. As such it is what I'll call a 'questioned source context'. In the second context, the subject has no such doubt or uncertainty so it isn't a questioned source context. In virtue of its not being a questioned source context it is what I'll call an 'unquestioned source context'. My proposal is that epistemic circularity in a questioned source context is malignant and that epistemic circularity in an unquestioned source context is benign (*op. cit.*, pp. 718–19).

Let me indicate very briefly some reservations about this approach.<sup>46</sup>

The view on offer is that bootstrapping produces a justified belief as to the reliability of a belief source *F* so long as there is no doubt or uncertainty that *F* is reliable (hereafter, '*RF*'). This claim raises a number of questions. What if you have creeping doubt about *RF*? Does bootstrapping give you less and less justification for that belief? Or do you encounter some special tipping point, where bootstrapping suddenly goes from providing full justification for *RF* to providing no justification for *RF* whatsoever? Further, imagine that there are anxiety and serenity pills people could take. The anxiety pills induce doubt about *RF*, the serenity pills dispel any such doubt. Does your taking the pills affect whether bootstrapping justifies *RF*? How could that be?<sup>47</sup>

These concerns about Bergmann's position presuppose that a subject's "having doubt" is a psychological notion. On this construal, a subject *S* has doubt with respect to *X* just in case *S* feels unsure about *X* or lacks confidence that *X*. The difficulties just pointed out might suggest that having doubt should be taken as an epistemic notion, instead.<sup>48</sup> Whether a subject has doubt would then be a matter of what evidence or justification the subject has. But there is some trouble in seeing how this idea would go. First, suppose that *S* has doubt with respect to *X* just in case *S* is not justified in believing *X*. Then, to say that you are justified in believing *RF* in "an unquestioned source context" (that is, one in which there is no doubt that *RF*) amounts to saying that you are justified in believing *RF* in a context in which you are justified in believing *RF*. This condition is completely empty, and does nothing to distinguish a range of situations in which bootstrapping is legitimate.

<sup>46</sup> Others have adopted views similar to Bergmann's, notably Paul Boghossian, "How Are Objective Epistemic Reasons Possible?" *Philosophical Studies* cvii (2001): 1–40.

<sup>47</sup> Again, I agree with Fumerton, "Epistemic Internalism."

<sup>48</sup> For classic treatments of doubt in the epistemic sense, see René Descartes, *Meditations*, in *The Philosophical Writings of Descartes*, Volume 2, John Cottingham, Robert Stoothoff, and Dugald Murdoch, trans. (New York: Cambridge, 1984), and J. L. Austin, "Other Minds," reprinted in Austin, *Philosophical Papers* (New York: Oxford, 1979), pp. 76–116.

Another proposal would be that *S* has doubt with respect to *X* just in case *S* has, on balance, some reason for believing the denial of *X*. In other words, bootstrapping is acceptable so long as there is, on balance, no evidence against *RF*. One difficulty with this suggestion is that it is hard to distinguish from the view that we have defeasible a priori justification for *RF*, regardless of bootstrapping.<sup>49</sup> In other words, what real epistemic contribution does bootstrapping make beyond whatever resources we command a priori, and why is it completely inefficacious in the face of any contrary evidence whatsoever? A further drawback is that our ordinary situation is one in which we have no evidence, on balance, that induction is unreliable. Hence, *LJI* should be perfectly in order, but that is very hard to swallow. Perhaps there is some other epistemic reading of “has doubt” that yields better results, but that has yet to be seen.

Bergmann’s account has a further dimension. Bootstrapping would not answer or satisfy a “skeptic” who doubted the reliability of *F*. However, your inability to demonstrate or convey your justification to a skeptic does not imply that you lack justification yourself.<sup>50</sup> Maybe so. My concern here has been with the capacity of bootstrapping to justify beliefs held by the person who carries it out. Whether bootstrapping has any value or use beyond that is another matter. Furthermore, the “skeptic” in this case would be someone who doubts the reliability of *F*. The presence of these doubts is, presumably, what prevents one’s bootstrapping from having any force for the skeptic. Such a claim appears to be no different, in substance, from the sort of view already considered, and the difficulties that have come to light remain. Neither a psychological nor an epistemic construal of doubt sustains the thesis that, in the absence of doubt, bootstrapping provides justification for believing that *F* is reliable.

<sup>49</sup> For a view akin to this, see Crispin Wright, “Some Reflections on the Acquisition of Warrant by Inference,” in Susana Nuccetelli, ed., *New Essays on Semantic Externalism, Skepticism and Self-Knowledge* (Cambridge: MIT, 2003), pp. 57–77.

<sup>50</sup> Bergmann (*op. cit.*, pp. 724–25).