



Reliabilism Leveled

Jonathan Vogel

The Journal of Philosophy, Vol. 97, No. 11. (Nov., 2000), pp. 602-623.

Stable URL:

<http://links.jstor.org/sici?sici=0022-362X%28200011%2997%3A11%3C602%3ARL%3E2.0.CO%3B2-G>

The Journal of Philosophy is currently published by Journal of Philosophy, Inc..

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/jphil.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The JSTOR Archive is a trusted digital repository providing for long-term preservation and access to leading academic journals and scholarly literature from around the world. The Archive is supported by libraries, scholarly societies, publishers, and foundations. It is an initiative of JSTOR, a not-for-profit organization with a mission to help the scholarly community take advantage of advances in technology. For more information regarding JSTOR, please contact support@jstor.org.

RELIABILISM LEVELED*

Max knows that the 6 o'clock train goes all the way to Montauk. He just asked at the information booth in the station, where the staffers are highly competent and have direct access to the master timetable. But suppose that, instead of asking at the booth, Max had consulted an old, out-of-date schedule, and that many of the routes and times listed there had changed. Even if the schedule happened to say that the 6 o'clock goes to Montauk, Max would not know that it does. Whether Max knows or not seems to turn on the fact that information provided at the booth is extremely reliable, while information provided by the out-of-date schedule is not. Accordingly, one might hold that knowledge just *is* a reliably true belief, or, alternatively, a belief that results from a process that reliably produces true beliefs.

To adopt such a view is to endorse reliabilism with respect to knowledge. Here, I distinguish two versions of this position, which I call *neighborhood reliabilism* and *counterfactual reliabilism*. I then raise the question of what reliabilists can and must say about reflective or higher-level knowledge (that is, knowledge about one's knowledge or beliefs). It will emerge that both versions of reliabilism encounter serious difficulties in this connection, and must be regarded as unsatisfactory. These results suggest that the traditional justification requirement for knowledge cannot be supplanted by a reliability requirement, as many now are inclined to suppose.

I. RELIABILITY AND RELIABILISM

Reliabilism about knowledge takes different forms. To some extent, one can distinguish "reliable belief" accounts of knowledge from "reliable process" accounts. Roughly, according to the former, whether you know that *P* depends upon whether your *belief* that *P* is prone to error. According to the latter, whether you know that *P* depends upon whether the *process* that produced your belief that *P* is prone to error. Reliabilism about knowledge differs from reliabilism

*I wish to thank Jonathan Adler, William Alston, Stewart Cohen, Peter Klein, Hilary Kornblith, and Chad Mohler for helpful comments and discussion. I am especially indebted to Richard Feldman. Parts of this paper have their origins in a contribution to a symposium at the American Philosophical Association's Central Division meeting in 1995; a more recent draft was read at the Franklin and Marshall Colloquium on Metaphysics and Epistemology in 1999, and I appreciate the helpful discussion on that occasion. The title and much of the impetus for this paper are due to the work of Alvin Goldman.

about justification, which seeks to explicate justification (rather than knowledge) in terms of reliability. The view I am considering involves no commitment to the notion that justification can be so understood.¹

A principal motivation for adopting reliabilism is the idea that truth is the ultimate epistemic norm, and that besides truth, there are no others. We want what we believe to be true. Moreover, because truth is so important, we want our beliefs to be not just adventitiously true, but *securely* true. Knowledge, then, is securely true belief, and knowledge is compromised to the extent that there is falsehood or the possibility of falsehood. The greatest security one could have would be certainty or infallibility. A belief would be certain if it could not possibly be wrong. That is, in any possible world in which you believe X , X is true:

- (1) Necessarily, $B(X)$ implies X .

and, equivalently:

- (2) Necessarily, $\neg X$ implies $\neg B(X)$.

Similarly, a belief-forming process would be infallible if it could not possibly go wrong, and this conception can be spelled out along the lines of (1) and (2).

But such logical invulnerability to error is often an unattainably high standard. Whenever someone believes a contingent proposition as the result of an inductive inference, it is logically possible for the belief she arrives at to be wrong. Hence, it seems that knowledge must demand something less than certainty or infallibility—that is, something less than truth in any logically possible situation. Knowledge may require reliability instead. What is crucial is that reliability, whether it pertains to beliefs or to processes, need not be total or absolute. Something may do Y reliably, even if it would fail to do Y under very extreme or extraordinary conditions. An alarm clock may be reliable, despite the fact that it would not ring if the power went out, or if the ceiling collapsed on it. Accordingly, reliability as it pertains to knowledge may be understood as truth in all situations that could arise, *except* for extraordinary or outlandish ones that we do

¹ Rather, the reliabilist I have in mind denies that there is any independent, substantive justification condition for knowledge; see below. It may well be, though, that the arguments I shall present also tell against reliabilism about justified belief. I should mention that there is a hybrid of justificationist and reliabilist accounts, sometimes called the “reliable indicator theory.” According to such a view, S knows that P only if S 's belief that P is based on evidence E , where E is a reliable indicator of P . I shall not discuss this approach here.

not, and need not, care about. Thus, perception, or particular beliefs formed by perception, may count as reliable, despite the fact that perception, or particular beliefs arrived at by perception, can go wrong under certain circumstances.

This understanding of epistemic reliability can be sharpened somewhat, as follows. Think of possible worlds as being farther away from the actual world to the extent that they differ from the actual world. A belief is reliable just in case it turns out to be true whenever it is held in a neighborhood N of worlds not too far away from the actual world; a process is reliable just in case it yields (mostly) true beliefs in a neighborhood N of worlds not too far away from the actual world. Proceeding in this vein, the reliabilist may set out as a condition of knowledge:

(3a) In N , $B(X)$ implies X .²

or:

(3b) In N , $B(X)$ by process P implies X .

or:

(3c) In N , all or nearly all beliefs $B(X_i)$ formed by process P are such that $B(X_i)$ implies X_i .

Assume, at least for now, that the boundary of N is fixed. In particular, it does not vary with the content of the proposition believed, the process of belief formation, or pragmatic considerations.

(3a)-(3c) invoke one conception of reliability, but there are others. And, clearly, there will be different versions of reliabilism corresponding to the different conceptions of reliability one might adopt. Theorists such as David Armstrong, Fred Dretske, Robert Nozick, and Keith De Rose have held that knowledge requires satisfaction of a co-variation or tracking condition:

(4) If X were false, then $\neg B(X)$.³

² Note that we do not want the converse condition for knowledge, namely, that in N , X implies $B(X)$.

³ Like (3), the tracking condition is subject to modification. (4) as stated corresponds to (3a) above. An analogue to (3b) would be:

(4b) If X were false, then $\neg(B(X)$ by process P).

where P is the process you actually used to arrive at your belief that X . Proponents of the tracking condition, like Nozick and De Rose, have in fact been led to include some reference to methods or processes in their accounts. But that has proved to be something of a quagmire; it is unclear how the relevant method or process is to be specified, and it is difficult to say exactly how reference to the process should figure in the

A belief might meet a version of condition (3) as stated, yet not meet condition (4). Here is why. Consider (3a) and (4). Given the standard Lewis-Stalnaker semantics for counterfactuals, instances of (4) are true if and only if the consequent is true in the closest possible world (or worlds) in which the antecedent holds. Which world that is, and how far it is from the actual world, depends upon what the value of X is. To illustrate, consider two propositions I know:

- (5) I am not carrying a pen with blue ink.
- (6) The Earth is not governed by cows.

Condition (4) makes my knowledge of (5) depend upon what I would believe in the closest possible world in which I was carrying a pen with blue ink, and it makes my knowledge of (6) depend upon what I would believe in the closest possible world in which the Earth was governed by cows. The closest $\neg(5)$ world—one in which I happen to be carrying a pen with blue ink—need not be very different from the actual world, in which I happen to be carrying a pen with black ink. That is, the closest $\neg(5)$ world is very close to the actual world, and we may assume that it falls within the boundary of N . A possible world in which the Earth was governed by cows would, however, be very different from the actual world, so different that it might well lie beyond the boundary of N . We may suppose that my belief as to whether the Earth is governed by cows is always true inside N , and that I thereby satisfy condition (3a). It remains open that I am wrong about this matter in the nearest world where cows do govern, contrary to (4). Thus, satisfaction of condition (3a) with respect to a given proposition does not guarantee satisfaction of condition (4) with respect to that proposition.⁴

In light of the foregoing, we can distinguish two different versions of reliabilism. The first, which I shall call *neighborhood reliabilism* (NR), is the view that you know X , if and only if you have a true belief that X which satisfies some version of (3). The second, which I shall call *counterfactual reliabilism* (CR), is the view that you know X , if and only if you have a true belief that X which satisfies some version of both (3) and (4).

It is important to note what requirements these theories do *not* impose on knowledge. First, NR and CR eschew the traditional claim

account. For further comment on (4b), see footnote 18. I cannot think of any version of the tracking condition that corresponds neatly to (3c).

⁴ Another way to make the point: suppose you know the proposition that the actual world is in N . Evaluation of whether that belief satisfies (4) requires considering the truth of what you believe in worlds outside N .

that one could not know a proposition without having justification for believing it. Let us say that to be justified in believing a proposition means something like having beliefs or nondoxastic states that provide good reasons for believing that proposition. Reliabilists who reject the justification condition sometimes argue as follows. The justification condition is independent of the reliability condition only if holding beliefs for good reasons is not necessarily truth conducive. In that case, the pursuit of rational belief and the pursuit of true belief could, at least in principle, diverge. But truth is the sole and ultimate epistemic value, so any substantive, independent justification condition has no place among the requirements for knowledge.⁵

That is one rationale for giving up the justification condition. A highly advertised benefit of doing so is that it immediately allows us to be rid of at least some forms of skepticism. According to a very familiar line of skeptical argument, I have no justification for believing that I am not a thoroughly deceived brain in a vat, and consequently I fail to know that I am not. Since I fail to know that I am not thoroughly deceived in this way, I have no knowledge of the external world. A reliabilist, or anyone else who rejects the justification condition, will balk at the first step. She may concede that I have no justification for believing that I am not a brain in a vat. Unless knowledge requires justification, however, it does not follow that I fail to know that I am not a brain in a vat, and the skeptical argument is blocked. I shall return to this issue later.

A second point is that reliability theorists typically reject various "higher-level" requirements on knowledge, such as:

(7) $K(X)$ entails $K(K(X))$.

and:

(8) $K(X)$ entails $K(R(X))$.⁶

The reliabilist has a number of motivations for denying (8), and, therefore, (7).

For one thing, a reliable belief or process is truth-tropic regardless of whether it is *known* to be so. Hence, a requirement that, beyond being reliable, a belief or process has to be known to be reliable

⁵ For an argument in this spirit, see Alvin Goldman, *Epistemology and Cognition* (Cambridge: Harvard, 1986), pp. 97-103.

⁶ Read ' $R(X)$ ' as ' S is reliable about X '. When I say ' S is reliable about X ', that locution is meant to stand in for whatever more specific reliabilist condition one might adopt. I shall assume, as a matter of terminology, that you are reliable about X only if you believe X and X is true. So, under this regimentation, $R(X)$ entails $B(X)$ and X . Moreover, for the reliabilist, $R(X)$ becomes equivalent to $K(X)$, and (7) becomes equivalent to (8).

would not do anything to foster the epistemic norm of truth. Such a requirement, then, has no place in an account of knowledge.

Next, from the reliabilist standpoint, you know X in virtue of being reliable about X . X is one fact about the world; whether you are reliable about X is another, different fact about the world. You are reliable about X if, in your environment, you are suitably related to X . You may be so related to X without being so related to $R(X)$. Hence, you can be reliable about X but not reliable about $R(X)$. You would, then, know X , but not know that you are reliable about X .

In addition, (8) as stated apparently requires an endless sequence of beliefs $B(R(X))$, $B(RR(X))$, $B(RRR(X))$, and so on, such that one is reliable about $R(X)$, $RR(X)$, $RRR(X)$, and so on. That is unattractive, certainly. And, if each of these beliefs has to be generated by distinct or additionally complex mechanisms or procedures, the regress is utterly vicious.⁷

Finally, the denial of (8) must be independently plausible, if reliabilism is to withstand one of the main criticisms which has been directed against it. This criticism, due to Laurence Bonjour⁸ and Keith Lehrer,⁹ is presented by way of various examples. Here is Bonjour's well-known "Norman" case:

Norman, under certain conditions which usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe the President is in New York City.... In fact the belief is true and results from his clairvoyant power under circumstances under which it is completely reliable (*op. cit.*, p. 62).

Bonjour maintains that Norman does not know in this instance, and concludes that reliabilism is untenable. Even among those who find this example convincing, there is some difference of opinion as to what defect in reliabilism it brings to light. But one prominent reaction is that Norman does not know the president is in New York because he fails to satisfy some higher-level requirement on knowledge.¹⁰ If Norman had verified that the deliverances of his clairvoyant

⁷ See Goldman, "Naturalistic Epistemology and Reliabilism," in Peter A. French, Theodore E. Uehling, and Howard K. Wettstein, eds., *Midwest Studies in Philosophy*, Volume XIX (Minneapolis: Minnesota UP, 1994), pp. 310-12. He claims that (7) and analogous principles create a vicious regress for nonreliabilist theories.

⁸ "Externalist Theories of Empirical Knowledge," in French, Uehling, and Wettstein, eds., *Midwest Studies in Philosophy*, Volume V (Minneapolis: Minnesota UP, 1980), pp. 55-73.

⁹ *Theory of Knowledge* (Boulder: Westview, 1990).

¹⁰ This is Bonjour's own conclusion (pp. 62-64; see also Lehrer, p. 165).

power reliably report how things are, then he would know that the president is in New York. But, since he lacks this higher-level knowledge, Norman lacks the first-level knowledge that the president is in New York.

I am not sure exactly what to make of this discussion, since I tend to lose my bearings when clairvoyance is so much as mentioned.¹¹ I do think, however, that this much is relatively clear: it is possible for one to be reliable about *X* without knowing that one is. According to reliabilist accounts of knowledge, simply being reliable about *X* is supposed to be sufficient for knowing that *X*. If knowing that you are reliable about *X* is an additional necessary condition for your knowing *X*, then these reliabilist accounts are mistaken.¹²

To the reliabilists I have in mind, the big picture looks like this. We know a lot about the world. We are reliable about things we believe on the basis of perception, so we know all kinds of things by seeing, hearing, and so forth. In addition, we are competent reasoners, at least about relatively elementary matters. We are reliable when we arrive at beliefs by induction or deduction, and these beliefs also count as knowledge. If we think about it, we probably believe that we are reliable about the things we believe on the basis of perception and inference. It is not apparent, however, whether these higher-level

¹¹ It seems to me, though, that a less far-fetched example to the same effect could be given. A few people have perfect pitch; they can immediately identify what absolute pitch they are hearing. I take it that whether someone does in fact have perfect pitch is not immediately apparent to that individual. To tell whether you have perfect pitch, you need to be tested. Let us suppose, then, that many people believe that they have perfect pitch, although they really do not. Now consider Norma, who has perfect pitch, but has never checked whether she does. She hears a tone, which she correctly takes to be an 'A'. Does she know that the tone is an 'A'?

¹² I should note that Goldman distinguishes "basic psychological processes" (my emphasis) of belief formation from *methods*, which are "various sorts of algorithms, heuristics, and learnable methodologies"—*Epistemology and Cognition*, p. 93. He also writes: "They are not part of the fixed, native architecture of the cognitive system. This is sufficient grounds for their not being basic processes" (p. 366). Goldman suggests that it is not enough for a method to be reliable to pass muster; it has to be acquired by a process or method that has higher-order reliability. That is, the first-order method needs to be acquired by a procedure that reliably leads to the acquisition of reliable methods (pp. 91-93). To this extent, a reliabilist account of knowledge modeled on Goldman's account of justification might recognize a higher-order requirement in the case of non-native, acquired belief-forming methods. But as Goldman himself acknowledges, the distinction between processes and methods is difficult to draw (p. 92 note 11). For example, the capacity to see requires maturation and stimulation in order to develop, but Goldman does not want perceptual beliefs to be subject to higher-level requirements. A further qualm is that it is hard to see how instituting such a requirement for methods can be squared with the thought that all we should ask for from beliefs or ways of forming beliefs is reliable truth.

judgments are themselves reliable. If and when they are not, then we fail to know that we know. But, says the reliabilist, there is nothing wrong with that.

II. COUNTERFACTUAL RELIABILISM IS TOO STRONG

I think the right picture is somewhat different. I agree that we know a lot about the world. I also agree that, when we form various beliefs, those beliefs or belief-forming processes may be reliable without our knowing that they are. Nevertheless, we often know that our beliefs about the world are true, not false. On some occasions we know, further, that we are reliable about what we believe. And, finally, there are many instances when we know that we know. It is incumbent upon reliabilists to get these facts about higher-level knowledge straight. I do not think they can.

Let me begin with CR. My treatment of it will be relatively brief, in part because I have discussed the status of (4) as a condition for knowledge elsewhere. (4) was the subject of intense scrutiny some years ago, and serious problems with it came to light. For one thing, it leads to unacceptable failures of what is known as the *closure principle for knowledge*:

(9) If $K(P) \ \& \ K(P \text{ entails } Q)$, then $K(Q)$.

Also, there are unexceptionable cases of inductive knowledge for which (4) does not hold, so (4) is too strong.¹³ There is an additional difficulty: (4) is hard to square with a satisfactory account of higher-level knowledge.¹⁴

Consider the following example. You see your long-time friend Omar, who is a perfectly decent and straightforward sort of person. Noticing his shiny white footwear, you say, "Nice shoes, Omar, are they new?" Omar replies, "Yes, I bought them yesterday." I think the following things are true:

(10) You know Omar has a new pair of shoes.

¹³ See my "Tracking, Closure, and Inductive Knowledge," in S. Luper-Foy, ed., *The Possibility of Knowledge: Nozick and His Critics* (Totowa, NJ: Rowman and Littlefield, 1987), pp. 197-215.

¹⁴ What follows is a more emphatic version of a criticism I originally directed against Nozick's account of knowledge (*ibid.*, p. 203). In his "Solving the Skeptical Problem," *The Philosophical Review*, civ (1995): 1-51, De Rose has defended something very like (4) (although, to be precise, De Rose's view is not that (4) is actually a condition for knowledge per se, but works instead to set the standards for correct knowledge attributions). De Rose takes note of the sort of problem I raised, for which he offers no solution (p. 22). He nevertheless retains confidence in a version of the counterfactual requirement, though on what basis I cannot see. For additional discussion of counterfactuals and knowledge, see my "Subjunctivitis," *Philosophical Studies* (forthcoming).

- (11) You know that your belief that Omar has a new pair of shoes is true, or at least not false.¹⁵
- (12) You know that your belief that Omar has a new pair of shoes is reliable.¹⁶
- (13) Other things being equal, you know that you know that Omar has a new pair of shoes.

(13) and (12) depend upon (11), and (11) is certainly compelling in its own right. (13) entails (11), because to know X entails that you have a true belief that X . (12) entails (11), if your being reliable about X is understood to entail that your belief that X is actually true.

The prime question at this point whether CR can accommodate (11). Does CR allow you to know that your belief that Omar has a new pair of shoes is true, not false? In particular, does (11) meet the condition set by (4)? I think the answer is no. I shall try to say why somewhat loosely and informally, and then more carefully.

(4) requires that, if you know X , then you would not believe X , if X were false. As things actually are, you believe that your belief that Omar has new shoes is not false. What if it were? If somehow your belief that Omar has a new pair of shoes were false, you would still believe that your belief was true, not false. The alternative is hard to fathom. It is difficult to conceive of your not believing that something you believe is true, whenever the matter happens to cross your mind. So, if your belief that Omar has new shoes were false, you would still believe that your belief was true, not false. You thereby fail to satisfy condition (4). According to CR, then, you do not know that your belief that Omar has new shoes is true, not false.

Let me now go through the same point a little more carefully. Let $O =$ 'Omar has new shoes'; $B(O) =$ 'You believe Omar has new shoes'; $B(\neg(B(O) \ \& \ \neg O)) =$ 'You believe that you do not believe falsely that Omar has new shoes'.

- (14) In the actual world, $B(\neg(B(O) \ \& \ \neg O))$.

Consider the nearest possible world W in which your higher-level belief is false, that is,

- (15) In W , $\neg(\neg(B(O) \ \& \ \neg O))$.

¹⁵ Note, in addition, that any proposition X itself entails that a belief that X is not false: X entails $\neg(B(X) \ \& \ \neg X)$. So, to say that you know that Omar has new shoes, but you fail to know that you do not believe falsely that Omar has new shoes would be to reject the closure principle for knowledge.

¹⁶ Alternatively: you know that the process by which you arrived at that belief is reliable.

Simplifying,

(16) In W , $((B(O) \ \& \ \neg O)$.

From (16),

(17) In W , $B(O)$.

If you believe O , you believe that you do not falsely believe O (see above). Hence,

(18) In W , $B(\neg(B(O) \ \& \ \neg O)$.

Given (15) and (18),

(19) In W , $(B(O) \ \& \ \neg O) \ \& \ B(\neg(B(O) \ \& \ \neg O))$.

That is, if your belief that O were false, you nevertheless would believe that it was not false. So, according to (4), you fail to know that you do not believe falsely that O . It follows, I take it, that you also fail to know that your belief is true.¹⁷

Note that the argumentation here is quite general. (4) makes it impossible for you to know that *any* of your beliefs is true, not false.¹⁸ Thus, it also seems to exclude your knowing that any of your beliefs is reliable, or that any of your beliefs is knowledge. That is going far too far, in my book.¹⁹ So, let us set CR aside, and move on to NR.²⁰

III. NEIGHBORHOOD RELIABILISM IS TOO WEAK

NR makes knowledge out to be true belief that satisfies some version of (3). The differences among the various formulations, and the

¹⁷ Unfortunately, this maneuvering is clumsy and roundabout. I did it this way in part to avoid complications about counterfactuals with disjunctive antecedents. Such technical hazards will arise any time one tries to apply (4) to knowledge of a conjunction. So far as I am concerned, that itself is good reason to be suspicious about (4).

¹⁸ Very much the same goes for the modification of (4) mentioned in footnote 3, namely, (4b). Suppose you believe O by some particular process P . It should be possible for you to know that you do not falsely believe O by process P . That knowledge requires $B(\neg(P(O) \ \& \ \neg O)$, where ' $P(O)$ ' now stands for the proposition that you believe O by way of process P . Suppose that your belief were false, that is, $(P(O) \ \& \ \neg O)$. Make the further, plausible assumption that your path from believing O by process P to a belief in the falsity of the conjunction would remain the same. Then, your belief in the falsity of the conjunction fails to satisfy condition (4b). In other words, you cannot know that you do not falsely believe O by process P .

¹⁹ This argument also cuts against reliabilism that accepts (4) without (3).

²⁰ A proponent of CR might attempt to surmount this difficulty by adding to her account a stipulation that you know any propositions you deduce from other propositions you know. This step is fraught with difficulties, however. The revised account may be too weak, and fail for the same reason that NR does (see below). Moving in this direction would also exacerbate other problems facing CR, such as those raised by Saul Kripke in widely heard but unpublished lectures.

forms of NR they give rise to, can be quite significant. The trend these days seems to be toward reliable-process theories, which run along the lines of (3c) rather than (3a) or (3b). Hence, in what follows, I shall consider the variant of NR that incorporates (3c), although what I shall have to say carries over to the other forms of NR as well.²¹ As I have set things up, NR is generally weaker than CR, since NR forgoes condition (4) while CR includes it. I have maintained that CR is too strong. It stints us higher-level knowledge we actually have. Now, I shall argue that NR is too weak. It would allow us higher-level knowledge we do not have.

Once again, I shall proceed by considering an example, namely, the “gas-gauge case,” due in its original form to Michael Williams.²² Williams describes himself driving a car with a working, highly reliable gas gauge. Williams does not know, however, that the gauge is reliable. Let us stipulate that he has never checked it, he has never been told anything about its reliability, and he does not even have any background information as to whether gauges like his are likely to be working. He never takes any special steps to see whether the gauge is going up or down when it ought to be. Rather, without giving the matter a second thought, Williams simply goes by what the gauge says. The gauge reads ‘F’, and Williams believes that his gas tank is full. According to NR, he *knows* that his tank is full. He has this knowledge because his belief results from a reliable process, that is, going by a well-functioning gas gauge. But Williams does not *know that he knows* that his tank is full. To have this higher-level knowledge, he would need to know that the gauge reliably registers the level of gas in his tank, and we have stipulated that he has no such information.²³

No one should be distracted by the details of the case. Perhaps it is hard to imagine how, in the situation described, Williams could fail to have *any* background information about the reliability of his gas gauge. Then again, one might doubt whether the supposedly reliable process Williams uses is appropriately identified as “going by a well-

²¹ For a somewhat different taxonomy of reliabilist positions, and for some arguments that would favor what he calls a “global” process theory like (3c), see Goldman, *Epistemology and Cognition*, pp. 43-49. I am not sure that Goldman’s official scheme has room for NR as I construe it, since he seems to assume that what he calls a “relevant alternatives” theory is a broadening of a “pure subjunctive theory” (p. 45). As I would put it, Goldman assumes the reliabilist accepts (4), and the open question is then whether to add some version of (3). I have set things up the other way around.

²² *Unnatural Doubts* (New York: Oxford, 1991), p. 347.

²³ Williams himself says: “No item of knowledge guarantees full knowledge of its own reliability conditions: indeed this is just another way of stating the externalist point that knowing does not guarantee knowing that one knows” (*ibid.*, p. 348).

functioning gas gauge." But what is essential here is just the claim that a subject might arrive at a belief by some reliable process, appropriately identified as *P*, yet not know that *P* is reliable. If particular features of my example create difficulties for you, let reading the gas gauge be a proxy for any such process *P*.

Now consider another driver, whom I shall call Roxanne. She is like Williams, in that she believes implicitly what her gas gauge says, without knowing that the gauge is reliable. But she undertakes the following, admittedly curious, procedure. She looks at the gauge often. Not only does she form a belief about how much gas is in the tank, but she also takes note of the state of the gauge itself. So, for example, when the gauge reads 'F', she believes that, on this occasion, the tank is full. She also believes that, on this occasion, the gauge reads 'F'. Moreover, Roxanne combines these beliefs; she believes:

(20) On this occasion, the gauge reads 'F' and *F*.

Certainly, the perceptual process by which Roxanne forms her belief that the gauge reads 'F' is a reliable one. By hypothesis, her belief that the tank is full is also reached by a reliable process. Hence, there seems to be no good reason to deny that Roxanne's belief in (20) is the result of a reliable process, and the reliabilist will say that she knows (20).²⁴

Now, it is a completely straightforward logical consequence of (20) that:

(21) On this occasion, the gauge is reading accurately.

Assume that Roxanne deduces (21) from (20). Deduction is certainly a reliable process, so there is no loss of reliability at this step. Consequently, it seems that Roxanne must be credited with knowing (21).

²⁴ What I have said may provoke disquiet of the following sort. You might believe that *A* by way of a reliable process, and believe that *B* by way of a reliable process, yet your belief that (*A* & *B*) might not be the result of a reliable process. For one thing, the chances of error with respect to *A* may be just barely acceptable and the chances of error with respect to *B* may be just barely acceptable. The chances of error with respect to the conjunction may then come out to be higher than the chances of error with respect to the conjuncts, and so the belief in the conjunction may not count as reliable. Another problem is that the process by which one comes to believe that *A* may interfere with the process by which one comes to believe that *B*, or vice versa. For example, one might use one instrument to measure for *A*, and another to measure for *B*, yet the operation of the second instrument may disrupt the functioning of the first. So, I think it is true in general that *R*(*X*) and *R*(*Y*) do not entail *R*(*X* & *Y*). But it seems implausible that such a failure occurs in the situation under discussion.

She knows this, supposedly, despite the fact that she has no independent information at all about the reliability of the gauge—whether it is broken or likely to be broken, whether it is hooked up properly, and so on. She just looks at the gauge and immediately believes what it says.

Let us say further that Roxanne does this over and over again for a good while. At various times t , she looks at the gauge, which reads 'X', and forms the belief that, on occasion t , the gauge says 'X' and the tank is X. Given what has already been said, she comes to know that, on each of these occasions, the gauge was reading accurately. Then, putting these pieces of information together, Roxanne concludes by induction:

(22) The gauge reads accurately all the time.

Reliabilists generally accept that induction is a reliable belief-forming process. So they should concede that the transition from various beliefs like (21) to (22) is knowledge preserving. We can add that, from (22), Roxanne infers:

(23) The gauge is reliable.

If Roxanne knows that the gauge is reliable, then she can presumably deduce and know that the process by which she comes to believe that her gas tank is full is a reliable one. And with just a little more deduction, Roxanne can come to know that she knows that her gas tank is full.²⁵

This extraordinary procedure, which I shall call *bootstrapping*, seems to allow one to promote many, if not all, of one's beliefs that were formed by reliable processes into *knowledge* that those beliefs were

²⁵ There is a noteworthy similarity between the problem I am raising for reliabilism and a difficulty that faces co-variational accounts of semantic content. The former is an attempt to analyze knowledge in terms of the co-occurrence of beliefs and the facts; the latter is an attempt to analyze semantic content in terms of the co-occurrence of linguistic or mental tokens and the facts. A difficulty for the latter is that, if tokens of R co-occur with distal events X , they also co-occur with the proximate process P that links X 's with R 's. So, the result would be that the R 's just as well represent P as X . See Fred Dretske, "Misrepresentation," reprinted in William Lycan, ed., *Mind and Cognition: A Reader* (New York: Oxford, 1992), pp. 129-43. The parallel problem I am raising is that, if the output of some belief-forming process P reliably co-occurs with X , that output also reliably co-occurs with instances of a reliable belief-forming process, namely, P itself. Thus, the result seems to be that P itself can give one knowledge that one is linked to X by way of a reliable process, that is, P can directly give one knowledge that one knows that X . The general lesson to be drawn is that co-occurrence can obtain for a variety of reasons, and is of questionable value in understanding relations like representation and knowledge.

formed by reliable processes.²⁶ I assume that bootstrapping is illegitimate. Roxanne cannot establish that her gas gauge is reliable by the peculiar reasoning I have just described. The challenge to NR is that it may go wrong here. On the face of things, it does improperly ratify bootstrapping as a way of gaining knowledge.²⁷

To sort this out, we need to get clear about what the defect in bootstrapping is. I shall consider some possible diagnoses, in a point-counterpoint format:

Point. It is obvious what Roxanne did wrong. Suppose the gauge had not been reliable. Roxanne would have gone through exactly the same procedure, and reached the false conclusion that her gauge was reliable. Hence, she does not know that the gauge is reliable.

Counterpoint. Say that if you like. But then you are really assuming that knowledge has to satisfy condition (4). In other words, you are abandoning NR in favor of CR, and we have already seen that CR fails.

Point. The problem with Roxanne's procedure is that it could not have possibly yielded any other result other than the one it did, namely, that the gauge is reliable.

Counterpoint. It is not clear that the putative defect really is one. The process by which I know I am conscious when I am is surely a reliable one, yet that process could not return a verdict other than that I am conscious.²⁸

Point. Roxanne does not know that the gauge is reliable, because she reaches that conclusion by bootstrapping, and bootstrapping itself is an unreliable process. After all, you can apply bootstrapping to a great many underlying processes, some reliable, some not. Every time, though, bootstrapping will tell you that the underlying process is reliable. When the underlying process is unreliable, bootstrapping will yield the false belief that the underlying process is reliable. So, bootstrapping itself often generates false beliefs, and must be considered unreliable.

²⁶ For definiteness, let us say that bootstrapping is the procedure that leads to beliefs, like (23), about the reliability of its underlying process. Bootstrapping may require that one be able to identify appropriately the underlying process by which one has arrived at a particular belief, and sometimes one may not meet that condition.

²⁷ Richard Fumerton independently arrived at a somewhat similar point—*Metaphilosophy and Skepticism* (Lanham, MD: Rowman and Littlefield, 1995), pp. 178–79.

²⁸ Something like this objection is raised by William Alston, *The Reliability of Sense Perception* (Ithaca: Cornell, 1993), p. 17. He does want to allow, however, that introspection counts as a reliable process (p. 139).

Counterpoint. So far as I can tell, this is the best answer for the reliabilist to give, but it still seems problematic in various respects. My first worry is a methodological one. The response under consideration identifies the process leading up to Roxanne's belief in (23) as a token of a relatively wide process type, roughly *bootstrapping in general*. The failures of other tokens of this broad type are supposed to discredit Roxanne's undertaking, and the belief that results from it. Such a response is dissatisfying because it is generally possible to find a process type under which a given process token falls, such that the token counts as reliable, just as it is generally possible to find another process type under which the token falls, such that the token counts as unreliable. There is no agreed-upon, principled way of identifying which is the proper process type to consider in evaluating the epistemic status of a particular belief.²⁹

The reliabilist herself may encounter some embarrassment on just this score. One might describe Roxanne's initially going by her gas gauge as "forming a belief by an *unauthenticated* process," that is, a process one does not know to be reliable. By parallel with the argument just given against bootstrapping, one might say that, if one forms beliefs by an unauthenticated process, often enough the token process one employs will be unreliable, and one's belief will be false. Hence, forming a belief by an unauthenticated process is generally unreliable, and beliefs that result from unauthenticated processes do not count as knowledge. Of course, the reliabilist rejects the demand that a process must be known to be reliable if it is to be a source of knowledge. She then has to find some basis for allowing the criticism of Roxanne's bootstrapping to stand, while rejecting the corresponding criticism of Roxanne's forming a belief by an unauthenticated process in the first place. Otherwise, the reliabilist's charge of guilt by association may boomerang.

Suppose we set this scruple aside, and accept that Roxanne's way of reaching (23) is unreliable. What makes it so? Where does the unreliability enter in? It will be useful to compare Roxanne's bootstrapping with a procedure that really would establish that a car's gas gauge is reliable. Imagine that Catherine periodically uses a dipstick to measure the level of gas in the tank of her car, and she verifies that

²⁹ This is the well-known "generality problem," which itself raises difficulties for reliabilism. See Richard Feldman, "Reliability and Justification," *The Monist*, LXVII (1985): 159-74; and Feldman and Earl Conee, "The Generality Problem for Reliabilism," *Philosophical Studies*, xciii (1997): 1-29. In presenting my own objections to reliabilism here, I have tried my best to steer clear of problems of this sort, allowing the reliability of the underlying process to be a matter of stipulation. But at this point such complications may be unavoidable.

the reading of the gauge matches that of the dipstick. Each time she does this, she comes to know an instance of:

(21*) On occasion t , the gauge is reading accurately.

After enough repeated successes, we can agree, Catherine may know by induction that the gauge is accurate all the time, and she can properly reach the yet stronger (modalized) conclusion that the gauge is reliable. The crucial point is that this impeccable way of finding out that a gas gauge is reliable is identical to Roxanne's original procedure from a certain point onward. Both Catherine and Roxanne infer that their gauges are reliable from beliefs of the form given by (21*). So, it seems that any lapse in reliability on Roxanne's part must occur in the way that she arrives at such beliefs.

In each case, Roxanne deduces an instance of (21*) from a belief of the form:

(20*) On occasion t , the gauge reads 'X' and X.

As I said earlier, it would be hard to deny that Roxanne's beliefs about how the gas gauge reads are reliably formed. And, by hypothesis, Roxanne's beliefs about how much gas is in her car's tank are formed by a reliable process. So, the reliabilist seems committed to the claim that when Roxanne comes to believe instances of (20*), those beliefs result from reliable processes. Given what has already been said, the only point left at which some defect can enter into Roxanne's procedure has to be in her transition from instances of (20*) to belief in corresponding instances of (21*). But, since any instance of (20*) entails an instance of (21*), that instance of (21*) must be true if the corresponding instance of (20*) is! In that sense, there is no more risk of error with respect to the instances of (21*) than there is with respect to the instances of (20*). So, it is opaque how proceeding from instances of (20*) to corresponding instances of (21*) turns a reliable process into an unreliable one.³⁰

³⁰ Williams insists that reliability cannot be lost through deduction (*op. cit.*, p. 328), leaving him little recourse at this point. Now, the thought may arise again that it is possible for processes to be reliable severally, but not in concatenation. One might then say that the process that led to Roxanne's belief in (20) is reliable, and that deduction is reliable, but that the combination of the two need not count as reliable. Goldman provides good reasons, however, for distinguishing processes that are reliable simpliciter (more precisely, "belief-independent J-rules") from those which are conditionally reliable ("belief-dependent J-rules"). The former generate (largely) true beliefs given nondoxastic inputs. The latter take beliefs as input, and will generate (largely) true beliefs, if their inputs are true. It seems to me that genuine deduction would be a paradigm of a conditionally reliable process, notwithstanding the logical errors made by subjects in certain experimental situa-

Let us also waive *this* objection for the time being, and permit the proponent of NR to take refuge in the claim that bootstrapping is an unreliable process. Thus, her position will be that you cannot use the beliefs generated by some process to establish the accuracy or the reliability of that very process. Unfortunately for the reliabilist, such a response to the gas-gauge case disrupts the full-blooded, noncontextualist antiscepticism that reliabilists often regard as a great virtue of their position. For, while bootstrapping may seem absurd when it involves someone's gas gauge, it has been regarded as a sound response to Cartesian skepticism.³¹

This answer to skepticism is basically G. E. Moore's vintage common-sense response, transposed and fortified by the addition of some early twenty-first-century reliabilism. It goes as follows. You know you have a hand. Your belief that you have a hand was formed by perception, which is a reliable process. Hence, that belief counts as knowledge. Now, the proposition that you have a hand entails that you are not a brain in a vat, deceived into believing falsely that you have a hand. Hence, you can deduce that you are not a brain in a vat from the proposition that you have a hand. Your belief that you are not a brain in a vat qualifies as knowledge, because it, too, is formed by a reliable process. At least, that will be so if deduction from a belief arrived at by a reliable process itself counts as a reliable process.

The correspondence between this neo-Moorean argument and the rejected line of argument in the gas-gauge case emerges clearly if the former is reformulated slightly:

- (A) You know you have a hand.
- (B) You know that it appears to you as though you have a hand.

tions. Goldman allows, as it seems he should, that a belief is reliably formed if it is the output of a conditionally reliable process whose inputs were the output of a reliable process or processes (*Epistemology and Cognition*, especially p. 83). Some care may be necessary here, but I think a reliable process theory must include some provision such as the one Goldman makes. Consequently, I have difficulty envisioning how the reliabilist can legitimately escape the difficulty raised in the text.

³¹ For an outline and discussion of reliabilism as an answer to skepticism, see Goldman, *Epistemology and Cognition*, especially pp. 55-57 (though he couches that discussion in terms of a reliabilist theory of justification). The bootstrapping aspect of such a response to skepticism emerges clearly in Alston, pp. 12-17, and in Williams, pp. 327ff. Alston endorses it reluctantly and with qualifications (pp. 138-40). Williams mistakenly believes that his own view avoids the kind of criticism I have raised. It is also noteworthy that reliabilists have embraced bootstrapping in mounting an inductive defense of induction. See, for example, James Van Cleve, "Reliabilism, Justification, and the Problem of Induction," in French, Uehling, and Wettstein, eds., *Midwest Studies in Philosophy*, Volume IX (Minneapolis: Minnesota UP, 1984), pp. 555-67; and David Papineau, "Reliabilism, Induction, and Scepticism," *The Philosophical Quarterly*, XLII (1992): 1-20.

(B) holds, because we may assume that your beliefs about how things appear to you are generated by a reliable process. Continuing:

- (C) Therefore, you know that your appearance of having a hand is veridical.
- (D) Therefore, you know you are not a deceived brain in a vat.

For comparison:

- (A') Roxanne knows that her gas tank is full.
- (B') Roxanne knows that the gas gauge reads 'F'.
- (C') Therefore, Roxanne knows that, on this occasion, her gas gauge is reading accurately.

Thus, the reliabilist version of Moore's refutation of skepticism sanctions the same kind of inference that created problems in connection with the gas-gauge case. Suppose we grant that the reliabilist is able to evade those problems by denying, either for principled reasons or by fiat, that bootstrapping can lead to knowledge. The upshot will be that reliabilism cannot provide a satisfactory response to skepticism along the lines just laid out.³²

Let me return to the status of NR in general. I have taxed that view for allowing bootstrapping to count as a source of knowledge, or for not being able to explain what is wrong with it. In the example given above, the crux of the matter came down to the status of beliefs of the form:

- (20*) On occasion *t*, the gauge reads 'X' and X.

and:

- (21*) On occasion *t*, the gauge is reading accurately.

It looked as though the reliabilist was committed to saying that Roxanne knows instances of (20*), but does not know corresponding instances of (21*), even though she deduces the second from the first. Now, contextualism about knowledge might appear to offer the reliabilist a way out of this uncomfortable situation. The contextualist will say that our standards for knowledge rise as we think through the example. That is, the neighborhood *N* which determines the reliabil-

³² One might think that the reliabilist can simply stand on the claim that we have perceptual knowledge of the external world, insofar as our perceptual beliefs are generated by a reliable process. We cannot, on this more minimal view, move on to knowledge that our experiences are veridical; that further step would allow for bootstrapping, which we have said is unreliable. One difficulty, among others, with such a position is that it is inconsistent with the closure principle for knowledge.

ity of processes expands. At first, when (20*) is under consideration, worlds in which the gas gauge fails are excluded from N . The gas gauge counts as reliable, and Roxanne is correctly judged to know both conjuncts of (20*), and (20*) itself. But overt consideration of (21*) brings questions about the accuracy of the gauge into the picture, and worlds in which the gauge is wrong come to be included in N . In this setting, the gauge no longer counts as reliable, so Roxanne does not know that her gas tank is full. Hence, she does not know one of the conjuncts of (20*), and she does not know (20*) as a whole. Nor does she know (21*), which she derived from (20*). Contextualism relieves the reliabilist's embarrassment of having to explain how reliability is lost in the course of Roxanne's deducing a proposition from another one that entails it.

I do not propose to enter into a full-blown discussion of contextualism here, but I shall make two brief points. First, if the standards of reliability and knowledge can be elevated in something like the manner just described, then we almost never know explicitly that we know anything. We also lose some seemingly impeccable first-order knowledge. To see this, recall the example in which you believe that Omar has new shoes. Say that you reflect upon your epistemic situation, and you believe that your belief that Omar has a new pair of shoes is the outcome of a reliable process. According to the contextualist revision of NR, considering the reliability of a belief-forming process expands the neighborhood N so that the process no longer counts as reliable. It follows that once you reflect upon the fact that your belief-forming process is reliable, you no longer know that Omar has a new pair of shoes. A fortiori, you do not know that you know that Omar has new shoes. Recourse to contextualism of this sort would make both first- and higher-order knowledge very fragile, unable to survive virtually any reflection at all. Such a result is hardly acceptable.³³

My second thought is that adding a contextualist dimension to NR may be going to great lengths to save a view that does not want to be saved. I say this because a consequence of the contextualist maneuver is that, if it is possible to maintain low epistemic standards while the entire bootstrapping procedure is carried out, that procedure would be successful. Roxanne could convert knowledge of how much fuel is in the tank of her car into knowledge that her gas gauge is reliable.

³³ To be sure, the contextualist scheme I have described is quite crude, and one might hope to find more satisfactory accounts of when and how epistemic standards get raised. But I am not aware of any worked-out proposal that would serve the reliabilist's purposes any better; see my "The New Relevant Alternatives Theory," *Philosophical Perspectives*, XIII (1999): 155-80, especially section 3.

But, for the reliabilist, knowledge of how much fuel is in the tank is knowledge of one empirical fact, and knowledge that the gauge is reliable is knowledge of another, independent empirical fact. One should not be able to transmute the first into the second, as bootstrapping would allow one to do.³⁴ To put the same point another way: suppose a change in epistemic standards of the sort envisioned does thwart or suppress the bootstrapping procedure. Still, the standards' tendency to change would not eliminate the possibility of bootstrapping in principle or somehow make bootstrapping intrinsically more acceptable. To that extent, whatever relief contextualism provides will strike the reliabilist as cosmetic, not real.

IV. SOME CONCLUSIONS

CR and NR, as I understand them, are meant to be general theories of what knowledge is. I have been arguing that both views encounter substantial difficulties when they have to deal with higher-level knowledge. CR allows:

(10) You know that Omar has new shoes.

but not the higher-level:

(11) You know that your belief that Omar has new shoes is true (not false).

NR got into trouble with the gas-gauge case. A proponent of NR has to say that Roxanne knows:

(20) On this occasion, the gauge says 'F' and *F*.

But if he goes on to grant that she knows:

(21) On this occasion, the gauge is reading accurately.

it appears that she can proceed to knowledge of:

(23) The gauge is reliable.

and beyond. This result is unacceptable, yet it is not clear how NR can avoid it.

What lessons should we draw? First, these phenomena provide some reason—not a conclusive reason, but some reason—to endorse the traditional view that knowledge requires justification. The traditional view handles the problem cases with relative ease. First, it offers a natural explanation of why (11) follows so closely on (10). Good evidence that

³⁴ I suppose Roxanne also has reliable beliefs, and knowledge, about the outputs of the gauge. What she still lacks is any access to whether the gauge itself is accurate or reliable.

Omar has new shoes is good evidence that a belief that Omar has new shoes is true, not false. If you have such evidence (and you believe that Omar has new shoes, and your belief is true, and there are no Gettier traps), then you are in a good position to have the knowledge reported in (10). Moreover, given that you have this evidence, you are also in a good position to have the higher-level knowledge reported in (11), so long as you recognize that you believe that Omar has new shoes, and you believe that your belief is true, not false.

Bringing justification into the picture also helps to explain what is wrong with Roxanne's bootstrapping. Given the way the example was framed, Roxanne has no justification for (23), nor, I think, for (22) or (21). If justification is necessary for knowledge, she does not know any of these, which seems correct. What about (20)? Since (20) directly entails (21), Roxanne would be justified in believing (21) if she is justified in believing (20). If, as I claim, she is not justified in believing (21), then she is not justified in believing (20), either. That, too, seems plausible. If Roxanne has no independent reason to believe that the position of the needle on the gauge is reliably correlated with how much gas is in the tank, I cannot see that she is justified in believing that her tank is full when she looks at the gauge. So the traditionalist can say that Roxanne does not know (20) or (21), since she lacks justification for both. There are no bootstraps for her to pull on.

Thus, re-instating a justification requirement for knowledge does make it easier to explain what is going on in the gas-gauge case. An objection is bound to arise, however. I have as much as said that for Roxanne to know how much gas she has by reading the gauge, she needs to have justification for thinking that the gauge is reliable. This claim seems to invite a version of the regress problem discussed in section 1. Suppose that, in general, justification for believing that X always requires justification for the belief that the process by which you came to believe that X is reliable. It apparently follows that justification for any belief will require an endless hierarchy of further justified beliefs of level $(N + 1)$ as to the reliability of the way one's belief at level N was formed.³⁵ But to say that justification *sometimes* requires reason to believe that one's belief-forming process is reliable, as it does in Roxanne's situation, does not imply that justification *always* requires such higher-level support or supplementation. It is this second, stronger claim that creates the regress, and one can refuse to agree to it. Note, though, that one can maintain that

³⁵ Some coherentists cheerfully accept this consequence.

knowledge always requires justification, without maintaining that justification always requires *further* justification for the higher-level belief that the process that led to the formation of one's initial belief was reliable.³⁶

In sum, reliabilism encounters serious difficulties in accounting for reflective knowledge, difficulties that might well be avoided by a justificationist approach to knowledge. The question now arises whether the weakness in reliabilism that has come to light is confined to its treatment of reflective knowledge, or whether there is some deeper, more far-reaching problem with reliabilism which vividly shows itself when reflective knowledge is under consideration. Correlatively, one might ask whether the only essential work the justification condition does is to capture the relations between lower- and higher-level knowledge.

My own view, which I shall state but not argue for, is that the inability of reliabilism to account for reflective knowledge has its roots in a more basic deficiency. I take it that knowledge is a kind of human success. It is something we accomplish; we have to do our part. The justification requirement is an attempt to formulate what that part is. Reliabilism either loses sight of this aspect of knowledge, or tries, in effect, to substitute some kind of *de facto* alignment between belief and the world for the knower's own contribution.³⁷ As a consequence, reliabilist theories go wrong in their treatment of various cases, and these include the examples deployed above. Reliabilism has trouble accounting for reflective knowledge because it is defective as an account of knowledge in general. A consideration of reflective knowledge shows that knowledge in general is neither reliably true belief nor belief that results from a reliable process.

JONATHAN VOGEL

Amherst College

³⁶ For such a view, see Alston, "What's Wrong with Immediate Knowledge?" reprinted in his *Epistemic Justification: Essays in the Theory of Knowledge* (Ithaca: Cornell, 1989), pp. 57-78. Another response would be to deny that the pertinent beliefs fall into the hierarchy of levels described in the text, setting off the regress. Perhaps first-order beliefs about how various mechanisms perform can justify beliefs about the reliability of those mechanisms, and about the reliability of belief-forming procedures that make use of those mechanisms. A view something like this was developed by Wilfrid Sellars, "Empiricism and the Philosophy of Mind," reprinted in *Science, Perception and Reality* (London: Routledge and Kegan Paul, 1963), pp. 127-96.

³⁷ So-called "virtue epistemology" seeks to do justice to the contribution of the knower while retaining many of the central features of reliabilism. My impression is that the reliabilist component of such accounts prevents them from giving a proper treatment of the way a knower's *evidence* figures in knowledge, but I shall not pursue that issue here.

LINKED CITATIONS

- Page 1 of 1 -



You have printed the following article:

Reliabilism Leveled

Jonathan Vogel

The Journal of Philosophy, Vol. 97, No. 11. (Nov., 2000), pp. 602-623.

Stable URL:

<http://links.jstor.org/sici?sici=0022-362X%28200011%2997%3A11%3C602%3ARL%3E2.0.CO%3B2-G>

This article references the following linked citations. If you are trying to access articles from an off-campus location, you may be required to first logon via your library web site to access JSTOR. Please visit your library's website or contact a librarian to learn about options for remote access to JSTOR.

[Footnotes]

¹⁴ **Solving the Skeptical Problem**

Keith DeRose

The Philosophical Review, Vol. 104, No. 1. (Jan., 1995), pp. 1-52.

Stable URL:

<http://links.jstor.org/sici?sici=0031-8108%28199501%29104%3A1%3C1%3ASTSP%3E2.0.CO%3B2-C>

³¹ **Reliabilism, Induction and Scepticism**

David Papineau

The Philosophical Quarterly, Vol. 42, No. 166. (Jan., 1992), pp. 1-20.

Stable URL:

<http://links.jstor.org/sici?sici=0031-8094%28199201%2942%3A166%3C1%3ARIAS%3E2.0.CO%3B2-3>