

## Externalism, Skepticism, and the Problem of Easy Knowledge

José L. Zalabardo

### 1. Bootstrapping

We think that some of our belief sources can furnish us with knowledge of the world—that I can come to know, for example, that the bottle is empty, by looking at it; that I had eggs for breakfast, by remembering it; that the Tory leader has resigned, by reading the newspaper; that my son has a temperature, by using a thermometer; or that 12 times 21 is 252, by mental arithmetic. Some of the most prominent debates in contemporary epistemology concern the conditions under which a belief source can be said to have this power.

One question that has attracted considerable attention in these debates is the status of the principle that in order for a belief source to have the power to produce knowledge, the subject needs to know that it is reliable. It will help matters to have a precise formulation of the principle. For this purpose we need to introduce a couple of concepts. On the standard tripartite conception of knowledge, which I shall assume in what follows, knowledge is a species of true belief: in order for a subject S to know that p, p has to be true, S has to believe that p, and S has to satisfy with respect to p an additional condition that turns her true belief into knowledge. I shall refer to this third condition as *warrant*.<sup>1</sup> Thus, when p is true and S believes that p, S will know that p just in case p has warrant for S. But if p is false or S doesn't believe that p, S will not know that p, even if p has warrant for S.<sup>2</sup> We can now use the notion of warrant to introduce the other concept that we need. Let's say that a belief source is a *knowledge source* (in circumstances C) just in case the fact that a belief has been formed as a result of its operation (in circumstances C) confers warrant on the belief.<sup>3</sup>

Now we can formulate our principle as a necessary condition for a belief source to be a knowledge source. For reasons that will become apparent later on, instead of discussing the principle that accords this status to the subject's knowledge of the reliability of the source, I shall concentrate on the weaker principle that substitutes warrant for knowledge:

WR A belief source is a knowledge source only if the proposition that the source is reliable has warrant for the subject.

For the last two decades or so, the dominant position on WR has been that the principle should be rejected. However, an unpalatable consequence of abandoning WR has been highlighted in the recent literature. It appears that once we give up WR, we are forced to condone an intuitively illegitimate method of knowledge acquisition.<sup>4</sup> In broad outline, the reasoning that yields this conclusion goes as follows. If we abandon WR, it becomes possible for a belief source to be a knowledge source even though the subject doesn't know it to be reliable. Consider Roxanne, who forms beliefs on whether the petrol tank in her car is full by checking whether the petrol gauge reads 'F'. Suppose that this belief source is a knowledge source, but the proposition that the source is reliable doesn't have warrant for Roxanne. Then it seems that Roxanne could mount an inductive argument for the reliability of her petrol gauge. Each of its premises would be of the form

Pi At time  $t_i$ , the gauge reads/doesn't read 'F' and the tank is/isn't full.

Roxanne's belief in each of these premises would be formed by looking at the gauge at time  $t_i$ . In light of our assumption, it follows from this that the second conjunct of each of these premises will have warrant for her. Hence, on the harmless assumption that Roxanne can come to know whether the gauge reads 'F' by looking at it, we can conclude that each of these premises has warrant for her.

Now, from P<sub>i</sub> Roxanne can infer, for every  $i$ :

A<sub>i</sub> At  $t_i$  the gauge is reading accurately.

And from a sufficiently broad and varied collection of premises of this form, she can derive inductively the following conclusion:

R The gauge is reliable.

Jonathan Vogel, whose presentation of the problem I have followed here, has referred to this procedure as *bootstrapping*.<sup>5</sup> Many epistemologists agree that once we accept that this belief source can be a knowledge source for Roxanne, even though she doesn't have warrant for the proposition that the source is reliable, we have to accept also that she can obtain warrant for this proposition by bootstrapping. And since rejecting WR allows this situation to arise, the rejection of WR is widely seen as carrying a commitment to the legitimacy of bootstrap-

ping as a method for acquiring warrant for the proposition that a belief source is reliable. It seems hard to deny that this is an undesirable outcome, and some have seen it as a *reductio* of any account of warrant that rejects WR.<sup>6</sup> But others have resisted this thought. For them, the acceptance of WR faces obstacles that should make us regard bootstrapping as the lesser evil.<sup>7</sup>

Perhaps the most influential line of reasoning against WR is the thought that rejecting the principle is the only effective response to a very powerful argument against the possibility of knowledge. The argument seeks to establish that our basic belief sources, such as sense perception, memory, introspection, and inductive and deductive reasoning, cannot be knowledge sources. And it takes WR as its starting point. Thus, for example, in the case of sense perception, WR entails that our perceptual beliefs would only have the status of knowledge if we knew that sense perception is reliable. But such knowledge, the argument goes, would require reasons for thinking that perceptual beliefs have a high true-to-false ratio, and such reasons would have to include propositions that we can only know by sense perception. It follows that we cannot have perceptual knowledge until we know that sense perception is reliable, and we cannot know that sense perception is reliable until we have perceptual knowledge. This circle, the argument concludes, renders perceptual knowledge impossible. And since the same reasoning can be applied to our remaining basic belief sources, this strategy can be used to challenge the vast majority of our knowledge claims. This piece of skeptical reasoning is often known as the *problem of the criterion*, as it was first used by Pyrrhonist philosophers to challenge the Stoics' claim to have found a criterion of truth.<sup>8</sup>

The threat of skepticism is not the only consideration blocking the adoption of WR. Other lines of thought should bear some of the blame for the unpopularity of the principle. Here I'd like to single out two of them. One is the impression that accepting WR would prevent us from ascribing knowledge in situations in which intuition dictates that knowledge is perfectly possible. According to this line of thought, there are cases in which we would want to say that the subject can acquire knowledge by applying a certain belief source, although she doesn't know that the source is reliable.<sup>9</sup> The other is the idea that the principle is incompatible with externalist and, in particular, reliabilist accounts of warrant.<sup>10</sup> If, as reliabilism dictates, reliability is a sufficient condition for a belief source to be a knowledge source, then a reliable

belief source will be a knowledge source, even if the subject doesn't know it to be reliable, contrary to what WR demands.

If these three thoughts are correct, WR makes knowledge impossible, precludes intuitive warrant ascriptions, and is incompatible with what many see nowadays as independently plausible accounts of warrant. In this situation, treating bootstrapping as a legitimate method of knowledge acquisition can easily appear as a small price to pay. My main goal in this paper is to defend WR against these three lines of reasoning. I shall argue in section 5 that blocking the skeptical argument doesn't require rejecting WR. Then, in section 6, I shall argue that WR doesn't have some of the counterintuitive consequences that have been attributed to it. Finally, in section 7, I question the claim that the principle is incompatible with reliabilist accounts of warrant. But before we turn to WR, I would like to consider, in sections 2 and 3, a very similar situation generated by another epistemological principle. The approach that I shall develop in section 3 for handling this case will provide the central ideas of my defense of WR.

## 2. A Similar Problem

The dialectical situation regarding WR can be seen in terms of a dilemma. On the one hand, we could accept the principle, but then we are in danger of finding a powerful skeptical argument irresistible, of having to accept important revisions to our intuitive knowledge ascriptions, and of being precluded from endorsing reliabilist accounts of warrant. On the other hand, we could reject the principle, but then we would have to accept an intuitively illegitimate method of knowledge acquisition.

Stewart Cohen has recently highlighted a parallel between the dilemma involving WR and a similar situation regarding the principle of closure under known entailment.<sup>11</sup> I believe the parallel is real and important, although, as we are about to see, the way I think it should be drawn differs from Cohen's approach in important respects. I shall start by presenting the parallel as I think it should be drawn and later consider how this differs from Cohen's presentation.

Let's start by formulating the principle that generates the dilemma in this case. As with WR, I am going to concentrate on a version of the principle that substitutes warrant for knowledge:

*Closure.* If  $p$  has warrant for  $S$ , and  $S$  knows that  $p$  entails  $q$ , then  $q$  has warrant for  $S$ .

The Closure Principle appears to generate a dilemma that mirrors closely the situation that we have described for WR. Notice, first, that the obstacles to the adoption of WR that we considered above have close correlates for the Closure Principle.

First of all, Closure, like WR, appears to play an essential role in a prominent skeptical argument. You know that your having a broken fingernail entails that you are not an incorporeal victim of an evil demon. Hence, by the Closure Principle, if you don't know that you are not an incorporeal victim of an evil demon, you don't know that you have a broken fingernail either. But you don't know that you are not an incorporeal victim of an evil demon—the argument goes. For adequate reasons for ruling out this possibility would have to include propositions that you can come to know only through procedures, such as sense perception, that would be reliable only if the evil-demon hypothesis were false. Hence using these procedures in this way would beg the question, as it would require presupposing that the evil-demon hypothesis is false. Therefore we can conclude that you don't know that you have a broken fingernail. The argument can be applied to every proposition that you know to be incompatible with a version of the evil-demon hypothesis that you don't know not to obtain, including the vast majority of the propositions for which we take ourselves to have warrant. Clearly, this argument rests on the Closure Principle, and some have advocated abandoning the principle as a strategy for resisting the skeptical conclusion that the argument seeks to establish.<sup>12</sup> Then admitting that we don't know that the evil-demon hypothesis doesn't obtain would not force us to conclude that we don't know any of the other things that we take ourselves to know.

The other two lines of thought against WR that I identified in the previous section also have close correlates for Closure. On the one hand, it has been argued that Closure is incompatible with our intuitions regarding the ascription of knowledge, since there are cases in which we want to ascribe to a subject knowledge of a proposition  $p$  that she knows to be incompatible with another proposition  $p^*$ , even though she doesn't know that  $p^*$  doesn't obtain. On the other hand, Closure has been attacked on the grounds that it is incompatible with a certain externalist account of warrant. As Robert Nozick argued, his tracking account of knowledge is incompatible with Closure, and consequently Closure would be undermined to the extent that the tracking account can be independently supported.

The parallels with the other horn of the dilemma are no less striking. As in the case of WR, abandoning Closure would seem to force us to accept an intuitively illegitimate method of knowledge acquisition. Consider one of the examples that Fred Dretske used to argue that Closure is in conflict with our intuitive knowledge ascriptions.<sup>13</sup> Assume that, as Dretske claims, by going to the zoo and looking at the zebra enclosure, you can come to know that the animals in it are zebras, even if you don't know that they are not mules cleverly disguised by the zoo authorities to look like zebras, but you know that this is entailed by their being zebras. Suppose a subject, call him Fred, finds himself in this situation, and after coming to know that the animals are zebras by looking at them, he goes on to derive the conclusion that the animals are not cleverly disguised mules from the premise that they are zebras. Now, since he knows the premise of this argument and he knows that it entails the conclusion, it seems hard to deny that by virtue of this inference Fred would come to know that the animals are not cleverly disguised mules. But this is counterintuitive. Intuition doesn't seem to treat Fred's inference as a legitimate method for acquiring this piece of knowledge.<sup>14</sup> This problem exhibits the same structure as the problem concerning bootstrapping. Following Cohen, I shall refer to both problems collectively as *the problem of easy knowledge*.

Thus, Closure and WR appear to generate dilemmas that exhibit important similarities. In what follows, I am going to exploit these similarities by arguing that a very natural strategy for dealing with the dilemma generated by Closure can be applied, with some modifications, to the dilemma generated by WR.

However, before we proceed, I want to discuss the differences between my presentation of the dilemma generated by Closure and Cohen's presentation. There is a crucial disparity between our accounts of how this version of the problem of easy knowledge arises. I have presented it as arising from the rejection of Closure, whereas Cohen blames the problem on the Closure Principle.<sup>15</sup> Let me explain why I think that Cohen's diagnosis is incorrect. This version of the problem of easy knowledge arises for a position that enables Fred to obtain warrant for the proposition that the animals are not cleverly disguised mules by deriving this conclusion from the premise that they are zebras, which in turn he has come to know by looking at them. But in order to make room for this maneuver, a position has to allow Fred to know that the animals are zebras and that this entails that they are not cleverly disguised mules, while not yet having warrant for the proposi-

tion that the animals are not cleverly disguised mules. This is the situation in which he'd have to find himself before performing the inference, if doing so is going to result in his obtaining warrant for the conclusion. Thus, far from being generated by Closure, the problem arises only for positions that reject this principle.

I shall not try to diagnose Cohen's mistake. I can suggest only that it might arise from a confusion between Closure and a principle along the following lines:

*Transmission:* If  $p$  entails  $q$  and  $S$  knows that  $p$ , then inferring  $q$  from  $p$  would enable  $S$  to obtain warrant for  $q$ .

But Transmission is different from, and independent of, Closure.<sup>16</sup> Transmission postulates a sufficient condition for an inference to have the power to furnish the subject with warrant for its conclusion. Closure, by contrast, imposes a constraint on an admissible combination of warrant attributions to a subject: that for all propositions  $p$ ,  $q$  and every subject  $S$ , if we ascribe to  $S$  knowledge that  $p$  entails  $q$ , we should not ascribe to her warrant for  $p$  without also ascribing to her warrant for  $q$ . Transmission may seem to sanction Fred's bootstrapping argument, but, as I have argued, it has this power only if Closure fails. It follows that those, like Cohen, who subscribe to Closure, have nothing to fear from this version of the easy knowledge problem. If Closure holds, the method of warrant acquisition licensed by Transmission will not, after all, enable Fred to come to obtain warrant for the proposition that the animals are not cleverly disguised mules by inferring it from the premise that they are zebras, given that he knows that the premise entails the conclusion. For if, on the one hand, he doesn't know the premise, then the antecedent of Transmission is not satisfied, and if, on the other hand, he does know the premise, then, by Closure, he already has warrant for the conclusion, and the inference would come too late to effect warrant acquisition.

### 3. The Closure Dilemma

I have argued that the second version of the problem of easy knowledge arises only if Closure is rejected, and hence that embracing Closure will make the problem disappear.<sup>17</sup> But a proponent of Closure has to find ways of dealing with the three lines of thought against the principle that I sketched above. My goal in this section is to present a proposal as to how to discharge these tasks.

Let's start with skepticism. If we accept Closure, my failure to know that I am not an incorporeal victim of an evil demon will entail that I don't know that I have a broken fingernail either. Hence the only way open for the proponent of Closure to resist the conclusion is to contest the skeptic's claim that I don't know that I am not an incorporeal victim of an evil demon. How does the skeptic defend this premise? In my informal presentation of the argument above, the claim that I don't know that I am not an incorporeal victim of an evil demon was supported on the grounds that my evidence for this proposition is bound to exhibit a feature that renders it inadequate. The precise construal of this disabling feature is a delicate issue. Broadly speaking, the idea is that my evidence for a proposition *p* will be inadequate if it is such that it would exhibit a certain shortcoming if *p* were false. The problem for the skeptic is to find a shortcoming that would be exhibited by my evidence for *p* if *p* were false and entails that my evidence is inadequate, even if *p* is true. Here I propose to waive these difficulties, since the strategy that I am going to present for dealing with the argument will not depend on whether the skeptic can find a satisfactory solution to this problem. I am going to construe the skeptical reasoning using one of the most plausible proposals as to which disabling feature invalidates my evidence against the evil-demon hypothesis—namely, that it relies on information that I can obtain only from sources that would be unreliable if I were an incorporeal victim of an evil demon. However, nothing will turn on this particular choice of feature.

On this construal, the skeptical reasoning will involve a premise asserting that my evidence is bound to exhibit the disabling feature:

- A All my evidence for the proposition that I am not an incorporeal victim of an evil demon relies on information that I can obtain only from sources that would be unreliable if I were an incorporeal victim of an evil demon.

It will also involve a premise asserting that the presence of this feature will invalidate my evidence:

- NC My evidence for a proposition *p* cannot be adequate if it relies on information that I can obtain only from sources that would be unreliable if *p* were false.

Clearly A and NC entail:

- B I don't have adequate evidence for the proposition that I am not an incorporeal victim of an evil demon.

But this is not quite the conclusion that the skeptic needs. The skeptic needs to establish the claim:

C I don't know that I am not an incorporeal victim of an evil demon.

But to progress from B to C, the skeptic needs to invoke an additional premise, which I shall label *Evidential Constraint on Knowledge* (EC):

EC In order to know that p, a subject needs to have adequate evidence for p.<sup>18</sup>

Without invoking the Evidential Constraint, the skeptic cannot derive C from B, and without C she cannot contrapose on the Closure Principle to mount a widespread challenge to our claims to knowledge. Even with Closure in place, the rejection of the Evidential Constraint would block the skeptical argument.

Hence the proponent of Closure has an alternative strategy for resisting the skeptical argument, provided that she is prepared to give up the Evidential Constraint. This is indeed a familiar strategy for dealing with this brand of skepticism—one that is commonly associated with externalism, since externalist accounts of knowledge typically license counterexamples to the Evidential Constraint. Vogel provides a very clear presentation of how the reliabilist could avail herself of this strategy (Vogel uses the term *justification* where I use *evidence*, and the label *Justification Condition* for the principle I'm calling *Evidential Constraint*):

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 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.<sup>19</sup>

These considerations suggest that the skeptical argument that invokes the Closure Principle can be resisted by advocates of the principle, provided that they are prepared to give up the Evidential Constraint. I shall not try to argue here that the Evidential Constraint should be rejected. Instead I am going to assume that the principle is false, and my defense of the approach that I am articulating will rest on this

assumption.<sup>20</sup> Hence my target audience will be restricted to those who are prepared to countenance the rejection of the Evidential Constraint. This still leaves me with a sizeable audience, since, as we have just seen, externalist epistemologists reject the constraint.

Let me now turn to the second line of reasoning against Closure that I sketched above, according to which the principle is in conflict with some of our intuitions concerning the ascription of knowledge. I want to suggest that the rejection of the Evidential Constraint also holds the key for dealing with the type of case that has been adduced in support of this objection to Closure. I think the point will become obvious if we look at the passage in which Dretske defends his claim that the zebra case constitutes a counterexample to Closure:

[T]hink a moment about what reasons you have, what evidence you can produce in favor of this claim. The evidence you *had* for thinking them zebras has been effectively neutralized, since it does not count toward their *not* being mules cleverly disguised to look like zebras. Have you checked with the zoo authorities? Did you examine the animals close enough to detect such a fraud? You might do this, of course, but in most cases you do nothing of the kind. You have some general uniformities on which you rely, regularities to which you give expression by such remarks as, “That isn’t very likely” or “Why should the zoo authorities do that?” Granted, the hypothesis (if we may call it that) is not very plausible, given what we know about people and zoos. But the question here is not whether this alternative is plausible, not whether it is more or less plausible than that there are real zebras in the pen, but whether *you know* that this alternative hypothesis is false. I don’t think you do. In this I agree with the skeptic.<sup>21</sup>

It is plain that Dretske’s case for the claim that you don’t know that the animals are not cleverly disguised mules is based on the contention that you don’t have adequate evidence for this proposition. Hence the argument will be blocked as soon as we abandon the Evidential Constraint. Then we can concede to Dretske that intuition dictates that you can know that the animals are zebras even though you have hardly any evidence for the proposition that they are not cleverly disguised mules, while refusing to treat the case as a counterexample to Closure. For in the absence of the Evidential Constraint, your lack of evidence is compatible with your knowing that they are not cleverly disguised mules.<sup>22</sup>

Let me now consider briefly the third argument against Closure, which seeks to undermine the principle on the grounds that it is incompatible with a tracking account of warrant. It seems to me that the incompatibility is undeniable, and hence that insofar as tracking

accounts are independently motivated, we will have reasons for abandoning the Closure Principle. Hence the approach that I am advocating will hold little appeal for those who are committed to tracking accounts. However, for those who are prepared to keep an open mind, incompatibility works both ways, and many nowadays have come to regard the independent plausibility of the Closure Principle as a powerful argument against tracking accounts. Notice that the connection between Closure and the problem of easy knowledge provides further support for this position. Given that tracking accounts are incompatible with Closure, and that the rejection of Closure generates the problem of easy knowledge, for proponents of tracking accounts the problem of easy knowledge is inescapable.

#### 4. Evidence and Warrant

Abandoning the Evidential Constraint amounts to rejecting a blanket demand for adequate evidence as a necessary condition for warrant. This rejection is perfectly compatible with the thought that there are circumstances in which warrant does require adequate evidence—or, more generally, that the amount and quality of the evidence that is required for warrant will vary from case to case. I shall not try to provide here a comprehensive account of the factors that determine how good our evidence for a proposition would have to be before it can be said that the proposition has warrant for us. I only want to single out one factor that seems to me to play a major role in fixing the evidence threshold for warrant. In the case of a true contingent proposition, the factor that I have in mind is how different things would have to be in order for the proposition to be false.<sup>23</sup> My claim is that the following principle will have to be an important ingredient of a correct account of how evidence is connected with warrant:

EW The quality of the evidence that a subject has to have for a true contingent proposition  $p$  in order to have warrant for  $p$  is, *ceteris paribus*, inversely proportional to how different things would have to be in order for  $p$  to be false.<sup>24</sup>

I shall not try to provide support for EW, although I don't think it would be hard to argue that our intuitive ascriptions of warrant follow this pattern.<sup>25</sup> I want to concentrate instead on how this principle, if correct, would affect our discussion of closure skepticism.

If we applied EW to Dretske's zebra example, it would explain why Fred needs very little by way of evidence for the proposition that the

animals in the zebra enclosure are not cleverly disguised mules in order to count as knowing this. This is so because things would have to be very different from what they actually are in order for the animals to be cleverly disguised mules. For this reason, the poor quality of Fred's evidence doesn't undermine the claim that he knows. Consequently, the claim that he knows that the animals are zebras cannot be attacked on the basis of his lack of adequate evidence, even with Closure in place.

If things didn't have to be so different in order for the proposition to be false, we would have a different situation. Suppose, for example, that Fred encounters the display labeled 'zebra enclosure', not in the zoo, but in an exhibition of British conceptual art. Then even if the zebra-looking animals he sees are real zebras, things wouldn't have to be all that different for them to be cleverly disguised mules. Hence, by EW, in order to know that the animals are not cleverly disguised mules, Fred would have to satisfy a higher threshold of evidence. If he didn't pass this test, he would not know that the animals are not cleverly disguised mules, and hence, by Closure, he wouldn't know that they are zebras either, even if, as it happens, that's what they are.

Notice that the limited extent to which the EW connects evidence with knowledge is much less than what the skeptic would need to rescue her argument. If the world is, by and large, as we take it to be, then my being an incorporeal victim of an evil demon would require things to be very different from how they are—vastly more different than they would have to be for the animals Fred is looking at to be mules. Hence, according to EW, if our conception of the world is broadly correct, knowing that I am not an incorporeal victim of an evil demon will not require having nonnegligible evidence for this proposition. I could know with little or no evidence, and hence the skeptical argument would still not go through.

Of course, our conception of the world could be radically wrong, and things might not need to be all that different in order for me to be an incorporeal victim of an evil demon.<sup>26</sup> In this case, EW will not protect my warrant for the proposition that this is not what I am from my lack of adequate evidence for it.<sup>27</sup> But I think we can concede this point to the skeptic without allowing her to rescue her argument. If the skeptic convinced me that my conception of the world is wrong in treating the evil demon hypothesis as far-fetched, I would conclude from this that I have no warrant for the proposition that I am not an incorporeal victim of an evil demon, and hence, by Closure, that I have no

warrant for any other proposition of which I know that it entails that the evil-demon hypothesis is false. But pursuing this line is well beyond the skeptic's resources. The skeptic never claimed to have an argument for the claim that the actual world is not radically different from how the evil-demon hypothesis portrays it—only an argument to show that I don't know that the hypothesis doesn't obtain. We are conceding that if she had the former she would have the latter, but since she doesn't claim to have the former, we are not conceding that she has the latter. Abandoning the Evidential Constraint defuses the skeptical argument, even after the link between evidence and warrant is partially restored by EW.

### **5. The Criterion**

With respect to the dilemma generated by Closure, I have advocated a position that preempts the problem of easy knowledge by endorsing the principle, and seeks to solve the ensuing skeptical difficulties by abandoning the Evidential Constraint. In this section, I am going to apply the same approach to the dilemma generated by WR.

Thus the first ingredient of the approach will be to deal with the problem of bootstrapping by endorsing WR. With WR in place, under no circumstances will a piece of bootstrapping allow a subject to acquire warrant for the proposition that one of her belief sources is reliable. For in order for an argument to have this effect, the subject's belief in each of its premises has to be warranted. If the argument is a piece of bootstrapping, this warrant will have resulted from the beliefs being formed with the belief source in question. But this requires that the belief source be a knowledge source, and this entails, by WR, that the proposition that the belief source is reliable has warrant for the subject. Therefore a necessary condition for a bootstrapping argument to confer warrant on its conclusion for the subject will be satisfied only if the conclusion already has warrant for the subject. Hence the argument will never allow the subject to obtain warrant for its conclusion.

We can illustrate the situation by looking at Roxanne's case. If, on the one hand, forming beliefs about the contents of her petrol tank by looking at the gauge is a knowledge source for her, then, by WR, the proposition that this belief source is reliable has warrant for her, and a bootstrapping argument can't result in her acquiring warrant for this proposition. If, on the other hand, this belief source is not a knowledge source, the premises of a bootstrapping argument for its reliability will

lack warrant, and hence, once again, the argument will fail to lead to the conclusion acquiring warrant for her.<sup>28</sup>

It should come as no surprise that my approach makes the problem of easy knowledge disappear, since it is universally accepted that the problem arises from the rejection of WR. What stands in the way of this solution are the formidable difficulties that we seem to confront if we accept WR. In the remainder of this section, I am going to focus on the first of the three problems that I sketched in section 1. I am going to articulate a proposal as to how to deal with the problem of the criterion without abandoning WR.

First of all, we need a more precise formulation of the skeptical reasoning. The argument seeks to undermine the status of a belief source as a knowledge source by challenging the possibility of warrant for the corresponding reliability claim. Hence WR is a crucial premise in the argument:

WR A belief source is a knowledge source only if the proposition that the source is reliable has warrant for the subject.

The possibility of warrant for reliability claims is questioned by reference to the evidence that we can obtain for them. Hence the argument needs to invoke, as in the closure-based skeptical reasoning, a principle connecting warrant and adequate evidence, but in this case the connection has to be assumed only for reliability claims:

ECR In order for a reliability claim to have warrant for a subject, she has to have adequate evidence for it.

To impugn our evidence for a reliability claim, the argument invokes the following principle:

NCR My evidence for the proposition that belief source S is reliable cannot be adequate if it relies on information that I can obtain only from S.<sup>29</sup>

These principles do not generate a problem when the reliability of a belief source can be supported by evidence that doesn't violate NCR—evidence that doesn't involve information that I can obtain only from the source in question. To complete the skeptical reasoning with respect to a specific belief source, we need to invoke an additional premise to the effect that its reliability cannot be supported by evidence that satisfies NCR. Thus, for example, in the case of sense perception, we would need the following principle:

SP All the evidence that I can gather for the reliability of sense perception relies on information that I can obtain only from sense perception.<sup>30</sup>

With these four premises in place, the conclusion that sense perception is not a knowledge source follows unproblematically.<sup>31</sup>

As I mentioned above, nowadays the dominant strategy for avoiding this conclusion is to abandon WR. This protects the status of a belief source as a knowledge source from any argument purporting to establish that the subject cannot obtain warrant for the proposition that the source is reliable. However, as we have seen, the price of this maneuver is that one must then treat bootstrapping as a legitimate method for obtaining warrant for this proposition.

The alternative strategy for dealing with the problem of the criterion that I am going to recommend consists in rejecting ECR.<sup>32</sup> It should be clear that if we abandoned ECR, the skeptical argument could no longer be run. Then we could concede to the skeptic that adequate evidence for the reliability of one of our basic belief sources is not to be had, but resist the further conclusion that the proposition that the source is reliable cannot have warrant for us. In this way, we could protect the status of our basic belief sources as knowledge sources from the skeptical reasoning without abandoning WR.<sup>33</sup>

But can ECR be abandoned? Obviously, abandoning ECR is not an option for advocates of EC, the principle that adequate evidence is a necessary condition for knowledge of *any* proposition, since ECR follows from EC as a special case. If ECR were endorsed only by advocates of EC, I could end my discussion of the problem of the criterion here, since I am working on the assumption that EC is false. However, this is not the situation. ECR has found support among some of the most eloquent critics of EC. The principle is explicitly endorsed by William Alston, in the course of his seminal defense of the legitimacy of bootstrapping as a procedure for acquiring knowledge of reliability principles. It has also been advocated by James Van Cleve, another early advocate of the legitimacy of bootstrapping.<sup>34</sup>

Both authors defend ECR on similar grounds. Alston presents his defense in the following passage:

What alternatives are there to justification by adequate reasons? Even if we are justified in accepting propositions that seem self-evident to us, that would not seem to apply here. Any tendency to suppose (I) [the claim that sense perception is reliable] to be self-evident can be put down to a confusion between self-evidence and being strongly inclined to accept the

proposition without question. It does not seem at all plausible to take it to be the sort of proposition that one cannot believe without its being true, or justified. Moreover, whereas propositions about one's own current conscious states, e.g., are plausibly regarded as "self-warranted," (I) is not plausibly so taken. Nor does it appear to enjoy any of the other forms of "direct" justification.<sup>35</sup>

A similar reasoning appears to be at work in Van Cleve's recent defense of ECR:

Principles affirming the reliability of our faculties are both *contingent* and *general*. For me, it is hard to see how a principle combining these features can be basic—at least if basic propositions comprise only those that are somehow immediately evident or obvious.<sup>36</sup>

Both authors explicitly accept the possibility of warrant in the absence of adequate reasons, thus distancing themselves from EC. Nevertheless, they refrain from taking the further step of rejecting ECR, on the grounds that the other ways in which a proposition might acquire warrant for a subject are not suitable in the specific case of reliability principles. They both consider similar alternatives—"self-evidence" and "self-warrant" in Alston's version of the argument, and "immediate evidence" or "obviousness" in Van Cleve's. They appear to have in mind the form of noninferential warrant contemplated by traditional foundationalism, to which I shall loosely refer as *givenist warrant*.<sup>37</sup>

I think we can agree with Alston and Van Cleve that even if one accepts the possibility of givenist warrant in general, it would be utterly implausible to suggest that reliability principles can obtain warrant for a subject in this way. But, like every argument by elimination, this one will work only if the options that are ruled out exhaust the range of possibilities. And at this point the Alston-Van Cleve argument for ECR is open to attack. Their reasoning could be resisted by any philosopher who accepts forms of warrant other than inferential and givenist warrant. This camp includes a large number of contemporary epistemologists, for whom a subject can come to have warrant for a proposition as a result of the instantiation of a natural property that constitutes a sufficient condition for warrant that can't be reduced to the inferential or givenist model.<sup>38</sup> Let me refer to this kind of warrant as *naturalistic warrant*, and to those who are prepared to countenance the possibility of naturalistic warrant as *naturalists*. I shall not try to argue here that there are alternatives to inferential and givenist warrant, or that naturalistic warrant should be counted among them. Instead, I shall adopt the same attitude as with regard to EC, introducing as an assumption of

my subsequent discussion the claim that naturalistic warrant is a genuine possibility. Hence for the remainder, my audience will be restricted to naturalistic epistemologists. My goal will be to convince them of the virtues of my approach. For the naturalist, ECR is not supported by the Alston-Van Cleve elimination argument, since it does not contemplate the possibility that reliability principles obtain warrant for a subject according to the naturalistic model.<sup>39</sup> Hence, unless an argument is provided to show that naturalistic warrant is not suitable for reliability principles, it is open to the naturalist to reject ECR, and hence to embrace WR without falling prey to the problem of the criterion.

## 6. Intuitions

Let me consider now the second family of obstacles to the adoption of WR. It has been argued that principles along the lines of WR should be rejected on the grounds that they are in conflict with our intuitions concerning warrant. There are cases, the thought goes, in which we would want to say that a belief source is a knowledge source even though the subject does not know that the source is reliable. Proponents of this line of thought have sought to elicit this intuition with two main kinds of case. In this section, I am going to argue that neither type of case can be used to undermine the position that I am advocating.

The first type of case that I want to consider involves cognizers that seem perfectly capable of acquiring knowledge with their belief sources but lack the conceptual sophistication that would be required for forming reliability beliefs. Nonhuman animals and human infants can be claimed to fall in this category, but even among adult humans the requisite conceptual sophistication is surely much less widespread than is sometimes assumed. Forming a reliability belief requires possessing concepts such as *belief source* and *reliability*, which cannot be expected to figure in most people's conceptual repertoires, but it would seem wrong to make knowledge the exclusive preserve of the intellectual elite who have these concepts at their disposal.<sup>40</sup>

It seems to me that these intuitions are perfectly sound. I want to say that by using his sensory devices, my one-year-old son can come to know that there is someone at the door, and I have no problem according the same treatment to my dog.<sup>41</sup> But my son, let alone my dog, can't form the belief that his sensory devices are reliable, for lack of the requisite concepts. It should be clear, however, that these intuitions don't put the slightest pressure on WR. They certainly undermine the

stronger principle that in order for a belief source to be a source of knowledge, the subject needs to *know* that the source is reliable. The reason is that, since knowing that *p* requires believing that *p*, the fact that neither my son nor my dog can believe that his sensory devices are reliable entails that they can't know it either. And then this principle would force us to accept that their sensory devices can't be knowledge sources. But none of this applies to WR. Having warrant for the proposition that *p*, unlike knowing that *p*, does not require believing that *p*. Hence the observation that my son and my dog can't form the belief that their sensory devices are reliable is perfectly compatible with the hypothesis that this proposition has warrant for them.<sup>42</sup> Our intuitions regarding these cases put no pressure on WR.

The second kind of case that I want to consider involves subjects that seem intuitively capable of acquiring knowledge with a belief source, even though they have no evidence for its reliability. This type of situation is clearly exemplified by the case that Vogel borrows from Michael Williams to introduce his discussion of bootstrapping:

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.<sup>43</sup>

Notice how Vogel supports the claim that Williams doesn't know that the gauge is reliable. He is drawing this conclusion from the stipulation that Williams has no evidence for the reliability of the gauge. This suggests that accepting the intuition that "going by what the gauge says" can be a knowledge source for Williams would not force us to abandon WR. All we would have to accept is that a belief source can be a knowledge source when the subject lacks adequate evidence for its reliability. But this admission would put pressure on WR only in the presence of ECR, and so far we haven't found any reason why the naturalist should subscribe to this principle.

Hence, of the two main types of case that have been invoked to exhibit a conflict between principles along the lines of WR and our intuitions, we have found that one poses no threat to WR, while to show that the other exhibits a conflict requires invoking a principle, ECR, which the naturalist can be expected to reject. Therefore, for the nat-

uralist, neither type of case reveals a conflict between WR and our intuitions.

### 7. WR and Reliabilism

I turn now to the third family of considerations against the adoption of WR that I want to discuss. So far I have argued that the naturalist should not be discouraged from adopting the principle either by the problem of the criterion or by our intuitions concerning when a belief source can be a knowledge source. It may seem, however, that for those naturalists who subscribe to a reliabilist account of warrant, embracing WR is still not an option, since WR appears to be incompatible with reliabilism. My goal in this section is to raise some questions for this line of thought.

Reliabilism, in its strongest form, postulates necessary and sufficient conditions for a belief to have warrant. It will be useful for our purposes to subdivide the issue of the compatibility of reliabilism with WR into two separate questions—namely, whether WR is incompatible with the sufficient conditions for warrant postulated by reliabilism and whether it is incompatible with the necessary conditions for warrant postulated by reliabilism.

Let me look first at the question concerning reliabilist sufficient conditions. We can frame the issue in terms of a simple logical point. The claim that S is a sufficient condition for a state of affairs X imposes very precise limits on what can count as a necessary condition for X—N can be a necessary condition for X only if S is sufficient for N.

Now, since WR expresses a necessary condition for a belief source to be a knowledge source, we can derive from this logical point a criterion for when a sufficient condition for a belief source to be a knowledge source is compatible with WR:

WR\* S can be a sufficient condition for a belief source to be a knowledge source only if S is a sufficient condition for the subject's having warrant for the proposition that the source is reliable.

The claim I am making is that a sufficient condition for a belief source to be a knowledge source is compatible with WR just in case it satisfies the criterion expressed by WR\*. This is hardly surprising, since WR\* is in fact equivalent to WR.

We can now use WR\* to assess the compatibility between reliabilism and WR. Clearly, a reliabilist sufficient condition for warrant can be

expected to entail a sufficient condition for a belief source to be a knowledge source. A reliabilist sufficient condition for warrant can be expressed in terms of the following schema:

RS If a belief has been formed with a belief source that is reliable and satisfies conditions C, then the belief has warrant.

It follows from a claim of this form that the reliability of a belief source and its satisfying the C-conditions are jointly sufficient for the belief source to be a knowledge source. Hence, in light of WR\*, it follows that a reliabilist sufficient condition for warrant will be compatible with WR just in case, for any belief source that satisfies the C-conditions, its reliability is sufficient for the proposition that the source is reliable to have warrant for the subject.

Clearly, any hopes that a reliabilist account of warrant will be compatible with WR will rest on what it treats as the C-conditions in the schema RS. To see this, notice that the unrestricted claim that a subject has warrant for the proposition that a belief source is reliable whenever the source is in fact reliable would be utterly counterintuitive.<sup>44</sup> I want to suggest, however, that when we look at the specific conditions under which reliabilists take reliable formation to be sufficient for warrant, the case for incompatibility starts to appear rather weak. I shall concentrate on the version of reliabilism put forward by Alvin Goldman.<sup>45</sup>

Central to Goldman's reliabilist account of knowledge and justification is a contrast between two types of belief sources: "basic psychological processes, on the one hand, and various sorts of algorithms, heuristics or learnable methodologies, on the other."<sup>46</sup> He refers to the former as processes and the latter as methods. Let's consider the situation concerning each of these types of belief source.

For beliefs that have been formed with a method, Goldman explicitly declares the reliability of the method insufficient for warrant. In order for such a belief to have warrant, in addition, the method needs to have been acquired (or sustained) by a reliable second-order process, that is, one that has a high ratio of reliable methods among the methods it generates.<sup>47</sup> Hence in the instance of RS that expresses Goldman's version of reliabilism, the C-conditions will include a clause to the effect that if the source is a method, it will have been acquired or sustained by a reliable second-order process.<sup>48</sup>

I want to suggest that this feature of the account significantly weakens the intuitive plausibility of the claim that reliabilism is incompatible with WR. The most obvious cases where we wouldn't want to say that

the reliability of a source confers warrant for the subject on the proposition that the source is reliable are cases in which the source has been more or less randomly adopted by the subject. The example that Goldman uses in this connection will illustrate the point:

[S]uppose our friend Humperdink has attended a series of talks on mathematics by a certain Elmer Fraud. These talks are not under the auspices of any certified educational institution, and Humperdink has been warned that Fraud has no credentials in mathematics. Humperdink hears Fraud enunciate numerous principles and algorithms, almost all of them defective. Nonetheless, being a complete novice—and a gullible one at that—Humperdink blindly accepts and applies them all. In one case, however, Fraud happens to teach a perfectly correct algorithm. Humperdink internalizes this one along with the others, and applies it to a relevant class of problems.<sup>49</sup>

In this kind of case, intuition is firmly set against the idea that the subject has warrant for the proposition that the belief source is reliable, even if the source is, as a matter of fact, reliable. Hence it would be hard to defend the compatibility of WR with a version of reliabilism that treated a belief source in these circumstances as a knowledge source. But Goldman's version of reliabilism doesn't face this problem. The second-order process requirement is introduced, precisely, to rule out these belief sources as knowledge sources. On Goldman's account, Humperdink's internalized correct algorithm is not a knowledge source, even though it is reliable, because it was acquired with an unreliable second-order process. Therefore, the intuition that the proposition that the algorithm is reliable has no warrant for Humperdink lends no support to the claim that Goldman's reliabilism is incompatible with WR. I shall not try to consider in detail how the point generalizes. I want only to suggest that when a belief source has been acquired with a reliable second-order process, the reliabilist should have no problem accepting the claim that the proposition that the source is reliable has warrant for the subject. And if this is so, the sufficient conditions postulated by Goldman's reliabilism for a method to count as a knowledge source are perfectly compatible with WR.

In the case of processes, the situation is very different. The additional conditions that a process has to satisfy in order to count as a knowledge source, on Goldman's account, are purely negative. So long as the subject's cognitive state at the time does not undermine the warrant she would otherwise have for the beliefs that she forms with a reliable process, the process will count as a knowledge source. This means

that a proponent of Goldman's reliabilism who wants to subscribe to WR would have to accept that in many cases the reliability of a process is a sufficient condition for the proposition that the process is reliable to have warrant for the subject.

I want to suggest that the reliabilist should have no qualms about accepting this point. Notice that the reliabilist is prepared to sanction the subject's use of a reliable process by ascribing warrant to the beliefs that she forms with it. For the reliabilist, using the process is the right thing for the subject to do, from an epistemic point of view. Now, in some cases the trust that the subject places in the process by using it will crystallize into a belief in its reliability. It seems to me that the considerations that support the epistemic sanction that the reliabilist is prepared to extend to the subject's use of her process would in principle also support, in these cases, ascribing warrant to her belief in its reliability. And if this suggestion is correct, the reliabilist sufficient conditions for a process to be a source of knowledge will also be compatible with WR.

Let me now move on to considering the compatibility of WR with the necessary conditions for warrant postulated by reliabilism. Here we seem to face a more serious problem. According to reliabilism, reliable formation is a necessary condition for a belief to have warrant. This would apply, in particular, to a belief in the reliability of a belief source—if it hasn't been formed with a reliable source, then it doesn't have warrant. But WR appears to license counterexamples to this principle. Suppose that S is a reliable belief source that is a knowledge source. Then, according to WR, the proposition that S is reliable will have warrant for the subject, whether or not she believes this proposition. And if she does believe it, her belief will have the status of knowledge, whether or not it has been formed with a reliable source. If her belief in the reliability of S has not been formed with a reliable source, we will have a violation of the necessary condition for warrant imposed by reliabilism.

I think that we have now identified a genuine source of conflict between WR and reliabilism—but one that the reliabilist could easily avert. All that would be required is a slight weakening of the sufficient conditions for warrant postulated by the account—namely a provision to the effect that a belief source will not count as a knowledge source if the subject believes it to be reliable and this belief has not been formed with a reliable source.<sup>50</sup> Notice, first, that if this provision is introduced the conflict disappears, as the resulting account would no

longer force the proponent of WR to ascribe warrant to reliability beliefs that the account is committed to treating as unwarranted. Furthermore, the introduction of the provision can be defended from principles with indisputable reliabilist credentials. If we are prepared, with Goldman, to withhold the status of knowledge source from a belief source that has been acquired with an unreliable second-order process, we should have no problem with according the same treatment to a belief source when the subject's belief in its reliability has been formed with an unreliable source.

Let me now consider briefly another objection often leveled against the principle that a belief source has to be known to be reliable in order to count as a knowledge source. The thought is that the principle generates an infinite regress.<sup>51</sup> The argument goes roughly as follows: if belief source S1 has to be known to be reliable in order to be a knowledge source, then the subject needs to have formed the belief that S1 is reliable with a belief source S2 that is a knowledge source. But this would require, in turn, that S2 is known to be reliable, and hence that the subject has formed the belief that S2 is reliable with a belief source S3 that is a knowledge source... I shall not try to provide a precise formulation of the argument, or consider how one might try to save from this attack the principle that a knowledge source has to be known to be reliable. I want to point out only that the objection doesn't apply to WR. Notice that WR doesn't require "an infinite hierarchy of reliability beliefs."<sup>52</sup> In fact, WR doesn't require any reliability beliefs at all in order for a belief source to be a knowledge source. Warrant for a reliability proposition is required, but a regress of warranted reliability propositions is not generated either. For the warrant of the reliability proposition does not arise from the operation of a knowledge source, but from the satisfaction of the condition that bestows on the belief source the status of knowledge source. Hence WR does not generate the need for further warranted propositions.<sup>53</sup>

### **8. WR, Externalism and Naturalism**

I would like to end with two brief remarks about how the claims that I have defended affect the standard conception of the range of epistemological options. My first point concerns the contrast between internalist and externalist views in epistemology. I am going to focus on Van Cleve's construal of the contrast, although many others would construe

it along the same lines. Van Cleve characterizes internalism in the following terms:

The core tenet of the internalist views is this: there is no first-order knowledge unless there is also higher-order knowledge with respect to the factors that make first-order knowledge possible.<sup>54</sup>

Externalist views are characterized accordingly by the rejection of this tenet. I argued in the preceding section that a minor modification of a standard version of reliabilism would render it compatible with WR. Hence it would be open to the reliabilist to adopt WR, thereby avoiding the problem of easy knowledge. The point I want to make is that this position would count as internalist in Van Cleve's taxonomy, since according to it, reliable belief formation makes knowledge possible only if there is higher-order knowledge (or at least higher-order warrant) with respect to this factor.<sup>55</sup> I am not going to discuss whether this is a satisfactory outcome. I simply want to point out that if this were correct, it would force us to abandon the standard conception of reliabilist accounts of warrant as fundamentally opposed to internalism. A reliabilist could embrace internalism without abandoning the central ingredients of her account of warrant.

The second point I want to make concerns the relationship between WR and naturalism. I think there is a widespread implicit assumption to the effect that acceptance of naturalistic warrant should be correlated with hostility towards higher-order requirements, such as WR—that naturalists should reject these requirements. In light of our discussion of WR, it appears that in this particular case, the assumption gets things precisely the wrong way round. I have argued that the standard obstacles to the adoption of WR can be regarded as genuine only by those who reject the possibility of naturalistic warrant. Hence, far from being a consequence of naturalism, the rejection of WR can be supported only from an antinaturalist point of view.

*University College London*

#### Notes

I am grateful to Mike Martin and an anonymous referee for the *Philosophical Review*. Their comments on an earlier draft led to substantial improvements. Thanks also to Mark Kalderon and Ross Ford.

<sup>1</sup> I am leaving Gettier counterexamples out of the picture for the sake of simplicity, since they won't play any role in my discussion. I am not taking sides in the debate between those who think that a suitable account of warrant should handle these cases and those who think that they should be dealt with

separately.

<sup>2</sup>On the possibility of warrant without belief, See Alvin Goldman's distinction between *ex ante* and *ex post* warrant; Alvin I. Goldman, "What Is Justified Belief?" in *Justification and Knowledge*, ed. George S. Pappas (Dordrecht: Reidel, 1979), 1–23.

<sup>3</sup>When S believes that p, I shall say that S's belief that p has warrant to express that p has warrant for S.

<sup>4</sup>William Alston was the first to point out that abandoning a principle along the lines of WR would have this kind of consequence. See William P. Alston, "Epistemic Circularity," *Philosophical Studies* 47 (1986): 1–28. Similar ideas were presented by James Van Cleve in his discussion of the Cartesian circle. See James Van Cleve, "Foundationalism, Epistemic Principles, and the Cartesian Circle," *Philosophical Review* 88 (1979): 55–91.

<sup>5</sup>See Jonathan Vogel, "Reliabilism Leveled," *Journal of Philosophy* 97 (2000): 602–23.

<sup>6</sup>For recent defenses of this view, see Richard Fumerton, *Metaepistemology and Skepticism* (Lanham, Md.: Rowman and Littlefield, 1995); Vogel, "Reliabilism Leveled"; Stewart Cohen, "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research* 65 (2002): 309–29.

<sup>7</sup>For a recent defense of this line, see James Van Cleve, "Is Knowledge Easy—or Impossible? Externalism as the Only Alternative to Skepticism," in *The Sceptics: Contemporary Essays*, ed. Steven Luper (Aldershot, Hampshire: Ashgate, 2003), 45–59.

<sup>8</sup>See Sextus Empiricus, *Outlines of Pyrrhonism*, trans. R. G. Bury (Buffalo, N.Y.: Prometheus Books, 1990), 2.85. Roderick Chisholm has used the same label to refer to a related but importantly different problem. See Roderick M. Chisholm, *The Problem of the Criterion* (Milwaukee: Marquette University Press, 1973).

<sup>9</sup>William Alston pioneered this line of attack against the principles that he called *higher-level requirements*. See, for example, William P. Alston, "What's Wrong with Immediate Knowledge?" *Synthese* 55 (1983): 73–95.

<sup>10</sup>The incompatibility of reliabilism with the principle that knowledge sources have to be known to be reliable is highlighted by Vogel and Van Cleve. See Vogel, "Reliabilism Leveled," 606–9, Van Cleve, "Is Knowledge Easy—or Impossible?" 46.

<sup>11</sup>See Cohen, "Basic Knowledge and the Problem of Easy Knowledge."

<sup>12</sup>This strategy was advocated by Fred Dretske and Robert Nozick. See Fred Dretske, "Epistemic Operators," *Journal of Philosophy* 67 (1970): 1007–23, and Robert Nozick, *Philosophical Explanations* (Cambridge: Harvard University Press, 1981).

<sup>13</sup>See Dretske, "Epistemic Operators."

<sup>14</sup>Of course, even though it would be, in principle, desirable to accommodate this intuition, as in the case of bootstrapping, we might have to reject it in the end in order to satisfy more important desiderata.

<sup>15</sup>See Cohen, "Basic Knowledge and the Problem of Easy Knowledge," 312–13.

<sup>16</sup>The distinction between Closure and transmission principles has been highlighted by Crispin Wright. See, for example, Crispin Wright, "(Anti-)Scep-

tics Simple and Subtle: G. E. Moore and John McDowell,” *Philosophy and Phenomenological Research* 65 (2002): 330–48, at 332.

<sup>17</sup> It might seem puzzling to claim that embracing Closure makes a version of the problem of easy knowledge disappear, since the adoption of Closure doesn’t make the acquisition of knowledge any harder. Notice however that the label ‘problem of easy knowledge’ is being used for the perceived need to condone an intuitively illegitimate method of warrant acquisition. Closure does solve this problem, even though this is not achieved by making warrant acquisition harder. Nevertheless, it would be wrong to claim that Closure makes knowledge easy. Closure dictates that we cannot ascribe warrant to Fred for the proposition that the animals are zebras without also ascribing warrant to him for the proposition that they are not cleverly disguised mules. It does not dictate that we have to ascribe warrant to Fred for either of these propositions. See note 28, where this point is developed in the case of bootstrapping.

<sup>18</sup> Evidential support is a matter of degree, and EC will be weaker or stronger depending on what level of support is required for evidence to count as adequate. Nevertheless, for our purposes, there is no need to be particularly specific on this point, since the argument for C doesn’t appear to trade on setting the threshold of evidence particularly high. If the constraint on adequate evidence imposed by NC is legitimate for some adequacy threshold, it will remain legitimate if the threshold is set very low.

<sup>19</sup> Vogel, “Reliabilism Leveled,” 606.

<sup>20</sup> I shall assume, that is, that every version of the principle is false—that for every nonnegligible degree of evidential support, it is possible in some cases to know that p, even if the subject doesn’t have evidence that supports p to that degree. See note 18.

<sup>21</sup> Dretske, “Epistemic Operators,” 1016.

<sup>22</sup> The effect of this move is to block an argument against the possibility of Fred knowing that the animals are not cleverly disguised mules. It would be a separate task to establish that Fred has this knowledge, that is, that there is a nonevidential sufficient condition for warrant that Fred satisfies with respect to this proposition. See my discussion of naturalism at the end of section 5, below.

<sup>23</sup> For a formal development of this notion, see David Lewis’s notion of *comparative possibility*. See David Lewis, *Counterfactuals* (Oxford: Basil Blackwell, 1973), 52–56.

<sup>24</sup> When p is false, the quality of the evidence required for warrant might be directly proportional to how different things would have to be in order for p to be true. However, since in this case it is not possible to know that p, our intuitions are naturally less firm than when p is true.

<sup>25</sup> See, in this connection, the principle that Steven Luper calls *variation thesis*. See Steven Luper, *The Sceptics: Contemporary Essays*, xx.

<sup>26</sup> If I am a victim of an evil demon, the quality of the evidence that I would need in order to have warrant for the proposition that I am not a victim of an evil demon will have to be determined in a different way. See note 24.

<sup>27</sup> EW makes the evidence threshold for warrant depend on how different things have to be in order for the relevant proposition to be false—not on the

subject's views on this matter.

<sup>28</sup>Notice that, while adopting this strategy will enable us to avoid treating bootstrapping as a legitimate method of knowledge acquisition, this is not achieved by imposing a constraint on how one can come to know that a belief source is reliable that bootstrapping violates. Hence, one could complain that the worry that a position makes knowledge of the reliability of belief sources too easy would not be solved by adopting WR, since the move wouldn't make such knowledge any harder. This point highlights the fact that there are two separate worries that we could have with respect to an account of warrant: one is that it renders bootstrapping legitimate, and the other is that it makes knowledge of the reliability of a belief source too easy. I am using the label *the problem of easy knowledge* for the first of these worries, and this problem would be solved by the adoption of WR. Clearly, though, adopting WR wouldn't by itself solve the second problem. The effect of the adoption of WR is that whenever your position treats a belief source as a knowledge source but fails to ascribe warrant to the subject for the proposition that the source is reliable, you have a conflict that needs to be resolved—either by stripping the belief source of the status of being a knowledge source or by ascribing warrant to the subject for the reliability claim. Whether knowledge of the reliability of a belief source will become harder as a result of adopting WR will depend on which of these two strategies is adopted in each case. Notice also that we might still need to impose constraints to rule out bootstrapping as a legitimate strategy that a subject could employ to bestow on a reliability claim an epistemic status other than warrant.

<sup>29</sup>Notice that NCR is entailed by NC.

<sup>30</sup>For a book-length defense of this claim, see William P. Alston, *The Reliability of Sense Perception* (Ithaca, N.Y.: Cornell University Press, 1993).

<sup>31</sup>As with EC, different versions of ECR will result depending on the degree of evidential support that we require for evidence to count as adequate. But the skeptical argument against the possibility of obtaining warrant for a reliability claim doesn't trade on setting the evidence threshold for warrant particularly high. If NCR imposes a legitimate constraint on adequate evidence for some level of evidential support, it will remain legitimate when the threshold is set very low.

<sup>32</sup>As with EC (see note 20), the proposal is to reject every version of the principle—even those that set very weak levels of evidential support as a necessary condition for knowledge.

<sup>33</sup>Notice that if we abandon WR, a transmission principle might force us to abandon NCR, but if we keep WR, NCR does not come under this kind of pressure. Hence if NCR is independently plausible (and, for what it's worth, I believe it is), this will be a further advantage of the approach to the problem of the criterion that I am recommending.

<sup>34</sup>Cohen has also endorsed the principle. See Cohen, "Basic Knowledge and the Problem of Easy Knowledge," 309.

<sup>35</sup>Alston, "Epistemic Circularity," 4–5.

<sup>36</sup>Van Cleve, "Is Knowledge Easy—or Impossible?" 51. This argument for ECR is Van Cleve's principal objection to a solution to the problem of the cri-

terion along the lines of the one that I am advocating, which he ascribes to Thomas Reid.

<sup>37</sup>I am borrowing the term from James Van Cleve, "Epistemic Supervenience and the Circle of Belief," *The Monist* 68 (1985): 90–104.

<sup>38</sup>One of the most influential examples of this position is the reliabilist view that, in some circumstances, if you have formed the belief that *p* with a reliable process, then you have warrant for *p*. See Alvin I. Goldman, *Epistemology and Cognition* (Cambridge: Harvard University Press, 1986), for a defense of this view.

<sup>39</sup>Alston's and Van Cleve's advocacy of this line of reasoning is extremely puzzling, since both have championed the view that givenist warrant is not the only alternative to inferential warrant. See, for example, Van Cleve, "Epistemic Supervenience and the Circle of Belief," 95–96; Alston, "Epistemic Circularity," 12–14.

<sup>40</sup>For an attack on this kind of elitism, see Alston, "What's Wrong with Immediate Knowledge?"

<sup>41</sup>See Fred Dretske, "Two Conceptions of Knowledge: Rational vs. Reliable Belief," *Grazer Philosophische Studien* 40 (1991): 15–30, for a compelling defense of these intuitions.

<sup>42</sup>Notice that even if my son (or my dog) doesn't believe that his perceptual devices are reliable, one could argue that his behavior manifests what we might call *subdoxastic acceptance* of this proposition: in using his perceptual devices to find out how things stand, he conducts himself as if he believed that they are reliable. To defend WR from the line of attack under consideration, it would suffice to argue that the range of propositions for which a subject can, in principle, have warrant includes those to which she displays this kind of subdoxastic acceptance, even if we concede that many other propositions are excluded from this range by the limitations of the subject's conceptual repertoire (it is hard to see, for example, under what circumstances my son could count, any time soon, as having warrant for the proposition that every even number is the sum of two primes).

<sup>43</sup>Vogel, "Reliabilism Leveled," 612. Williams presents the example in Michael Williams, *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism* (Oxford: Blackwell, 1991), 347ff.

<sup>44</sup>For a counterexample to the unrestricted claim, see the Humberdink case that Goldman describes in the passage quoted below.

<sup>45</sup>As presented in Goldman, *Epistemology and Cognition*. See also Alvin I. Goldman, "Strong and Weak Justification," in *Epistemology*, Philosophical Perspectives 2, ed. James E. Tomberlin (Atascadero, Calif.: Ridgeview, 1988), 51–69.

<sup>46</sup>Goldman, *Epistemology and Cognition*, 93.

<sup>47</sup>See *ibid.*, 51–53, 93–95.

<sup>48</sup>In *Epistemology and Cognition* there is no explicit discussion of the possibility of higher-order methods, beyond a passing allusion to going beyond second-order processes (see p. 53). However, there is no obvious reason why Goldman shouldn't make room for this possibility. This would require a slightly more sophisticated account of how a first-order method would have to be acquired in order to qualify as a knowledge source.

<sup>49</sup> Goldman, *Epistemology and Cognition*, 52–53.

<sup>50</sup> Notice that this provision would not treat reliably formed belief in a belief source's reliability as a necessary condition for this belief source to count as a knowledge source. If the subject didn't believe the source to be reliable, the provision wouldn't prevent the source from counting as a knowledge source. The provision is activated only when the subject believes the source to be reliable *and* this belief has not been reliably formed.

<sup>51</sup> See Alvin I. Goldman, "Naturalistic Epistemology and Reliabilism," in *Philosophical Naturalism*, Midwest Studies in Philosophy 19, ed. Peter French, Theodor Uehling Jr., and Howard K. Wettstein (Notre Dame: University of Notre Dame Press, 1994), 301–20, at 311ff.; Vogel, "Reliabilism Leveled," 606–7.

<sup>52</sup> Goldman, "Naturalistic Epistemology and Reliabilism," 311.

<sup>53</sup> The version of the difficulty based on supervenience that Goldman considers doesn't apply either. As Goldman himself notes, "Perhaps supervenience could still be satisfied, for example, if (first-order) instantiation of the CRP property [the reliabilist sufficient condition for justification] itself necessitated satisfaction of clause (B) [the demand for justified belief in the reliability of the process]" (ibid., 312). Goldman's objection to this point is that "no such necessitation obtains" (ibid.). I think he is right in this, but notice that his argument for this claim rests crucially on the fact that (B) involves the subject having higher-order beliefs. Hence, the objection doesn't carry over to my proposal, and supervenience will be satisfied by virtue of the fact that the higher-order requirement imposed by WR (warrant for the proposition that the source is reliable) is necessitated by the naturalistic sufficient condition for the belief source to be a knowledge source.

<sup>54</sup> Van Cleve, "Is Knowledge Easy—or Impossible?" 45.

<sup>55</sup> In the formulation that I quote, internalism calls for higher-order knowledge, but Van Cleve makes it clear later on that a view that demanded higher-order warrant would also count as internalist. See *ibid.*, 45–46.