Understanding admits of degrees. So our response to testimony need not and often should not be either unqualified acceptance or unqualified rejection. Testimony is typically accompanied by tacit riders, caveats that as everyone in the relevant community knows restrict the scope or strength of the claims. We cannot understand testimony or incorporate its deliverances into our developing systems of thought unless we are sensitive to its tacit riders. We need not, of course, accept them, any more than we need accept the proffered statement of fact. Our community may add or delete riders of its own. But to give testimony its epistemic due, we must recognize how the considerations that go without saying circumscribe and inform the message that testimony transmits.[22]

[22] I am grateful to Jonathan Adler, Robert Fogelin, Jerrold Katz, Mary Kate McGowan, Amelie Rorty, and two anonymous referees for help with this paper.

Basic Knowledge and the Problem of Easy Knowledge

STEWART COHEN
Arizona State University

I Basic Knowledge and the KR Principle

It is a truism that we acquire knowledge of the world through belief sources like sense perception, memory, and induction. But when one thinks about how such knowledge is structured, one encounters a version of the problem of the criterion. For a natural intuition (pretheoretically anyway) is that a potential knowledge source, e.g., sense perception, can not deliver knowledge unless we know the source is reliable. But surely our knowledge that sense perception is reliable will be based on knowledge we have about the workings of the world. And surely that knowledge will be acquired, in part, by sense perception. So it looks as if we are in the impossible situation of needing sensory knowledge prior to acquiring it. Similar considerations apply to other sources of knowledge like memory and induction. Skepticism threatens.

The dominant response to this problem of the criterion focuses on the alleged requirement that we need to know a belief source is reliable in order for us to acquire knowledge by that source. Let us call this requirement, "The KR principle":

> KR A potential knowledge source K can yield knowledge for S, only if S knows K is reliable.

One can avoid the problem of the criterion by rejecting KR. Precisely how one rejects this principle will depend on one's other theoretical commitments.

Evidentialism holds that empirical knowledge has its source in evidence. The KR principle as interpreted by evidentialist views says that in order for S to know some empirical proposition P, S must possess some evidence E, and S must know that E is a reliable indication of P, or perhaps that E makes P probable.[1] An evidentialist who denies the KR principle holds that one can

[1] Versions of this principle have been proposed by Bonjour, Fumerton, and Sosa (forthcoming).

know P on the basis of evidence E without knowing that E is a reliable indication of P.[2] For example, one can know that X is red, on the basis of its looking red, without knowing that X's looking red is a reliable indication of X's being red.

Certain reliabilist views deny that knowledge, at least in some instances, has its source in evidence, at least as traditionally conceived. Rather a belief P can be knowledge in virtue of some 'external' relation it bears to the world, e.g., being reliably connected to the fact that P, or being produced by a reliable cognitive faculty. Such theories also reject the KR principle. For one's belief P can be knowledge provided the source of one's belief is reliable, even if one does not know that the source is reliable.[3]

Theories that reject the KR principle escape the problem of the criterion by allowing that a belief source can deliver knowledge prior to one's knowing that the source is reliable. Let's call such knowledge, "basic knowledge".[4] Is it plausible to hold that there can be basic knowledge? Here intuitions might diverge. Many philosophers have no intuitive qualms about basic knowledge whatsoever. And those that might have intuitive reservations—How can I know P on the basis of S if I don't know that S is reliable?- -have learned to live with basic knowledge, mainly because they can see no other way to avoid the problem of the criterion.

Now presumably, we do know our belief sources (or sustainers) are in some important sense reliable and so this is something that any theory of knowledge must account for. The problem with the KR principle is that it says that we can have no knowledge prior to our knowing that our belief sources are reliable. This makes it very difficult to account for how we in fact know that our belief sources are reliable—and so, given the KR principle, very difficult to account for any knowledge.

Theories that allow for basic knowledge have a distinct advantage here, since they can appeal to our basic knowledge in order to explain how we know our belief sources are reliable. According to such views, we first acquire a rich stock of basic knowledge about the world. Such knowledge, once obtained, enables us to learn how we are situated in the world, and so to learn, among other things, that our belief sources are reliable.

Consider as a representative of such a view, Ernest Sosa's virtue epistemology. According to Sosa, if one has epistemically virtuous faculties—faculties that produce, apt, reliable, beliefs—then one can come to know things via

> ...perception and introspection, along with intuition, as well as inductive and abductive reasoning, along with ...deductive reasoning... By use of such faculties...one attains... a broad view of oneself and one's environing world. And, if all goes well, then in terms of this epistemic perspective one can feel confident about the reliability of one's full complement of one's faculties.[5]

Sosa also notes that

> ...reflective knowledge, while building on animal knowledge [secured by reliable, apt faculties] goes beyond it precisely in the respect of integrating one's beliefs into a more coherent framework. This it does especially through attaining an epistemic perspective within which the object-level animal beliefs may be seen as reliably based.[6]

Alvin Goldman argues (roughly) that we can know simply in virtue of the reliability of our cognitive processes. Goldman's view thus allows for basic knowledge. Discussing metaknowledge Goldman says:

> To know that we...know, we would have to know that we...use reliable processes. But since the analysis makes it (logically) possible for us to know what processes we use, and makes it (logically) possible for us to know all sorts of truths about the world (which is essential for knowing the reliability of our processes), the analysis makes it possible for us to have higher-order knowledge.[7]

Let us say that views of this kind, i.e., views which hold that reliability knowledge is based on basic knowledge, have a basic knowledge structure (BKS).[8] What I want to argue today is that BKS views, by rejecting the KR principle, face serious difficulties. Once one accepts basic knowledge, we can get a plausible enough story about knowledge of the reliability of our knowledge sources. According to BKS views, such knowledge is attainable, but only after we acquire a substantial amount of information about the world. The problem is that once we allow for basic knowledge, we can acquire reliability knowledge very easily—in fact, all too easily, from an intuitive perspective. This suggests that we were wrong to think we had the basic knowledge in the first place. And this is just what the KR principle says. We can call this "The Problem of Easy Knowledge".

But if we do accept the KR principle, then we are faced with the problem of the criterion. So I will explore some possibilities for avoiding the problem without denying the KR principle.

---

2 Klein, Pollock (1986), Pryor.

3 Alston, Dretske, Goldman, Nozick, Sosa (forthcoming).

4 My notion of basic knowledge is not exactly the same as the traditional foundationalist notion. One could consistently hold that basic knowledge in my sense must be based on other beliefs, provided it need not be based on the belief that the belief source is reliable. So a coherentist theory could accommodate basic knowledge, in my sense. But of course any basic knowledge in the traditional foundationalist sense will be basic in my sense as well.

5 Sosa (forthcoming).

6 Sosa (forthcoming).

7 Goldman.

8 See Van Cleve (1979) and (1984).

II The Problem of Easy Knowledge: Closure

The easy knowledge problem arises in two related ways. The first derives from the way in which basic knowledge interacts with the principle that knowledge is closed under known entailment:

> If S knows P and S knows P entails Q, then S knows (or at least is in a position to know) Q.

This principle seems to me to be something like an axiom about knowledge. Notoriously, Nozick has used his tracking view of knowledge to deny the closure principle.[9] But most philosophers are unwilling to go along with Nozick here. They take the failure of closure as a strong reason to reject the tracking view. Moreover, as Kripke has demonstrated in his (regrettably unpublished) lectures on Nozick, the tracking view leads to instances of closure failure that even detractors of the closure principle can not endorse.

Let me first illustrate the problem closure presents for basic knowledge by looking at an evidentialist view—a view we can call evidentialist foundationalism. (Later we will look at non-evidentialist views.) According to this kind of a view, I can know, e.g., that the table before me is red, merely on the basis of its looking red (to me). On some versions, I know the table is red on the basis of my believing it looks red. But given that we do not obviously have beliefs about how things appear to us, it is perhaps more plausible to hold that I know it's red on the basis merely of its looking red. This view allows for basic knowledge in my sense because I can know the table is red, prior to knowing that the table's looking red is a reliable indication of its being red

This view can look intuitively quite acceptable. Provided I have no positive reason to believe X's looking red is *not* a reliable indication of X's being red, then perhaps I can know X is red on the basis merely of its looking red.[10] The knowledge thus acquired is relatively modest. Of course this example will be an instance of a more general principle that one can know X is ϕ on the basis of its looking ϕ, without knowing that ϕ-looking things tend to be ϕ. For this principle to work, there must be some restriction put on the sort of properties that ϕ ranges over.[11] And this can present a problem for explaining how we are able to acquire enough basic knowledge to eventually learn facts about our perceptual process, e.g., that such processes are reliable. But I do not wish to highlight the difficulty this view has in explaining how one ascends from one's stock of basic perceptual knowledge to knowledge about one's perceptual faculties.

On the contrary, the problems I will discuss reveal that it is all too easy to get from one's stock of basic knowledge to knowledge about one's perceptual faculties, given the closure principle. Because if I know the table is red, then it follows, by closure, e.g., that I know that I am not a brain-in-a-vat being deceived into falsely believing that the table is red. Now some proponents of evidentialism, possibly following some remarks of Moore, openly embrace this result as a way of responding to the skeptical argument that relies on the closure principle.[12] They concede that I can not come to know I am not a brain-in-a-vat being deceived into thinking the table is red, by inferring it directly from the table looks red. But I *can* come to know the table is red, on the basis of its looking red. And once I know the table is red, I can infer and thereby come to know that I am not a brain-in-a-vat being deceived into thinking the table is red.

Is this a plausible view? I think it may look plausible only because it is obscure in general how we know global skeptical alternatives do not obtain, e.g., how we know we're not brains-in-a-vat. And so insofar as we are inclined to say we do know such things, this can seem like a reasonable hypothesis about how we know.

But the problem is that we cannot limit the knowledge we can acquire in this way to denying global skeptical alternatives. For example, if I know the table is red on the basis of its looking red, then it follows by the closure principle that I can know that it's not the case that the table is white but illuminated by red lights. Presumably, I cannot know that it's not the case that the table is white but illuminated by red lights, on the basis of the table's looking red. So the evidentialist foundationalist will have to treat this case analogously to the global deception case: I can know the table is red on the basis of its looking red, and once I know the table is red, I can infer and come to know that it is not white but illuminated by red lights. But, it seems very implausible to say I could in this way come to know that I'm not seeing a white table illuminated by red lights.[13] Note that on this view, my inductive evidence against the possibility that there are red lights shining on the table turns out to be irrelevant to my knowing the table is not white with red lights shining on it. This is surely a strange result.[14]

The defender of evidentialist foundationalism might say that we must distinguish between knowing the falsity of the *alternative*, the table is white

9 Nozick.

10 If I know surely there must be some connection between the table's looking red and its being red, I just don't need to know there is such a connection. I will fail to know the table is red if I have some reason to believe that its looking red is not connected to its being red. This will defeat its looking red as a reason to believe it is red.

11 Pollock (1986) restricts the properties to those expressed by ostensive concepts and Pryor restricts the properties to those that are represented by the perceptual state.

12 Klein, Pollock (in conversation), Pryor.

13 I raise this problem in Cohen (1988) and (1999).

14 This view allows that inductive evidence can be relevant to defeater defeaters, e.g., if presented with the defeater that Jones says that the table has red lights shining on it, inductive evidence that Jones is not reliable could defeat this defeater. But when there are no defeaters, it will not be relevant.

but illuminated by red lights and knowing the falsity of the *defeater*, the table is illuminated by red lights. The falsity of the alternative is entailed by *the table is red* and so given the closure principle, it follows from one's knowing that the table is red that one knows the alternative is false. The falsity of the defeater is not entailed by *the table is red* and so it will follow that one knows it is false only if we accept the defeater elimination principle:

> If S knows P on the basis of R, and D is a defeater of R as a reason to believe P, then S knows D is false.

Since the defender of evidentialist foundationalism can accept the closure principle without accepting the stronger defeater elimination principle, s/he is committed only to our knowing the falsity of the alternative. Moreover one can infer the falsity of the alternative from *the table is red* using little more than disjunctive addition. For *it's not the case that the table is not red but illuminated by red lights* is equivalent to *the table is red or the table is not illuminated by red lights.*

I, myself don't find this response very convincing. It's counterintuitive to say we could in this way know the falsity of even the *alternative* that the table is white but illuminated by red lights. Suppose my son wants to buy a red table for his room. We go in the store and I say, "That table is red. I'll buy it for you." Having inherited his father's obsessive personality, he worries, "Daddy, what if it's white with red lights shining on it?" I reply, "Don't worry—you see, it looks red, so it is red, so it's not white but illuminated by red lights." Surely he should not be satisfied with this response. Moreover I don't think it would help to add, "Now I'm not claiming that there are no red lights shining on the table, all I'm claiming is that the table is not white with red lights shining on it". But if evidentialist foundationalism is correct, there is no basis for criticizing the reasoning.

Now it's true that one can infer that the table is not white but illuminated by red lights from *the table is red*, using little more than disjunctive addition. But the trivial entailment is not what's in question here. And neither is the closure principle which allows you to know the table is not white with red light shining on it on the basis of your knowing it's red. Both sides of the dispute agree that *if* we know the table is red, *then* we know it's not white but illuminated by red lights. The point of the objection is that it seems implausible that we could come to know such a thing in this way. And if we can not know it is not white but illuminated by red lights in this way, then given the trivial entailment and the closure principle, it follows that we can't know that it is red simply on the basis of its looking red.[15]

[15] Contextualism can mitigate somewhat the force of this objection. But it will remain true that the sentence "S knows it's red on the basis of it's looking red, and on the basis of it's looking red S knows (or at least has sufficient evidence to know) it is not white with red lights shining on it" is true in an everyday context.

I originally presented this as an objection to evidentialist foundationalist views. But it's easy to see that this version of the easy knowledge problem arises as well for non-evidentialist views like reliabilism. Now, as we noted, reliabilist theories can be viewed as strategies for avoiding the problem of the criterion. Reliabilism responds to this problem by jettisoning the requirement that you need to know a process is reliable in order to gain knowledge by that process—in effect denying KR. So consider perception. According to Reliabilism, I do not need to know my perceptual processes are reliable in order for me to look at the table and thereby come to know it is red. It's enough that my perceptual processes are reliable. But this leads to the problem of easy knowledge. For if one can know just by looking at the table that it is red, then one can infer and thereby come to know that it is not white but illuminated by red lights. So imagine again my 7 year old son asks me, "Daddy but what if it the (seemingly red) table is white with red lights shining on it?"... I reply, "Well—look, the table is red." According to reliabilism, this is something I can know provided my vision is reliable. So I can appeal to it in reasoning. So I continue. "But since it's red, it can't be white with red lights shining on it. See?" I take it that this reasoning is no more acceptable in this case then it was when we were considering evidentialist foundationalism.

Now one kind of non-evidentialist view can escape this result, viz., the tracking view endorsed by Dretske and Nozick.[16] On this view, even though I can know the table is red just by looking, I won't be able to know the table is not white with red lights shining on it. This is why the tracking view denies the closure principle. But again, I take this to be an unacceptable result. Imagine again my son worrying about whether the table is really red or just white with red lights shining on it. This time I claim to *know* the table is red. So my son asks me how I can be sure—Do I know it's not white with red lights shining on it? I reply, "Well, I don't know *that*, but that's an entirely separate matter from whether I know it's red." Of course, this dialogue exhibits the seeming incoherence regarding metaknowledge that results from the denial of closure.

The problem I have been discussing arises from denying the KR principle whether the theory that denies it is evidentialist or non-evidentialist. If you allow for basic knowledge, there is nothing to stop us from acquiring, by trivial inferences, all sorts of knowledge about how we are not deceived or misled by our belief sources.[17] But intuitively this knowledge seems avail-

[16] Dretske, Nozick.

[17] Is this knowledge of the reliability of our faculties? Not in the modal sense. It's evidence that this particular exercise of one's faculty is not an instance of deception.

able only if we have prior knowledge about the world—knowledge not required by the views in question.

III The Problem of Easy Knowledge: Bootstrapping

The second way in which the easy knowledge problem arises has been brought to our attention by Jonathan Vogel and Richard Fumerton.[18] Both Vogel and Fumerton raise the problem as an objection to Reliabilism, though I will later argue that, in fact, the problem generalizes to any theory that denies the KR principle. As applied to Reliabilism, the problem arises as follows: Suppose I have reliable color vision. Then I can come to know, e.g., that the table is red, even though I do not know that my color vision is reliable. But then I can note that my belief that the table is red was produced by my color vision. Combining this knowledge with my knowledge that the table is red, I can infer that in this instance, my color vision worked correctly. By repeating this process enough times, I would seem to be able to amass considerable evidence that my color vision is reliable, enough for me to come to know my color vision is reliable.

Vogel calls this process "bootstrapping". Clearly I cannot use bootstrapping to acquire knowledge that my color vision is reliable. But the reliabilist appears unable to explain why I can not.

Now there are various responses the reliabilist might make. If one accepts the tracking version of Reliabilism, then one can avoid this untoward result. Since bootstrapping would lead me to believe my color vision is reliable, even if it were not, it follows by tracking that I can not use bootstrapping to come to know my color vision is reliable. But I am proceeding on the supposition that the tracking view is incorrect.[19]

[18] Fumerton, Vogel (2000).

[19] A second response for the reliabilist, noted by Vogel (2000), is that bootstrapping is itself an unreliable process. For if you apply bootstrapping to an unreliable process P, it will still produce the false belief that P is reliable.

As Vogel notes this kind of response raises the issue of what is to count as a process. Moreover there doesn't seem to be any point in the bootstrapping process that the reliabilist can identify as where the unreliability enters in. Finally Vogel argues that for the reliabilist to argue that BS is unreliable would undermine the reliabilist's favored way of responding to skeptical arguments.

Aside from Vogel's arguments, I don't think that reliabilist can use the alleged unreliability of bootstrapping to avoid this version of the easy knowledge problem. To see why, we need to pay attention to a distinction introduced by Goldman in his pioneering 1979 paper—the distinction between conditionally reliable processes and unconditionally reliable processes (belief dependent and belief independent processes?). Goldman notes that some processes, e.g., inference, are belief dependent—they take beliefs as inputs. And the most we can require of these processes is that they produce true beliefs as outputs given true beliefs as inputs. i.e., that they be conditionally reliable. Since bootstrapping is a belief dependent process, i.e., it takes beliefs as inputs, the most we can require of it is that it be conditionally reliable. And, of course, bootstrapping is conditionally reliable, mainly because bootstrapping insofar as it is belief-dependent involves enumerative induction, and enumerative induction is conditionally reliable.

One might be tempted to raise other objections about why bootstrapping cannot produce *knowledge* of the reliability of a faculty. Perhaps the sort of evidence bootstrapping produces, what Alston calls "track-record evidence" (p1 & Bp1, p2 & Bp2,...,pn & Bpn) is insufficient to generate knowledge of reliability, the latter being a modal notion.[20] But this would be to miss the fundamental problem. Let's back up a bit from full-fledged bootstrapping. Consider just the single case where I infer from P and the fact that my believing P was produced by color vision, that color vision has generated a true belief—in this instance alone. Or more generally, the single case where I infer from P and the fact that I believe P that the process that produced P has generated a true belief. Surely the fact that the process produced a true belief this time is *some* evidence for the reliability of the process. And this is true whether or not by amassing further evidence of this kind, one could eventually come to *know* that the process is reliable. Yet surely it is strongly counterintuitive to say that one can in this way acquire any evidence for the reliability of one's process. *Ex hypothesi* one need not have had such evidence *prior* to the exercise of the process. According to Reliabilism, if the process that produces one's belief in P is reliable, then one can thereby know P, even where one has no evidence for the reliability of the process. But then surely one does not acquire such evidence simply by noting that in addition to P, one believes P, and drawing the inference that the process responsible for P has produced a true belief.

Imagine again my 7 year old son asking me if my color-vision is reliable. I say, "Let's check it out." I set up a slide show in which the screen will change colors every few seconds. I observe, "The screen is red and I believe it's red. Got it right that time. Now it's blue and, *look at that*, I believe its blue. Two for two..." I trust that no one thinks that whereas I previously did not have any evidence for the reliability of my color vision, I am now actually acquiring evidence for the reliability of my color vision. But if Reliabilism were true, that's exactly what my situation would be. We can call this the problem of "easy evidence"

Both Fumerton and Vogel raise the possibility of bootstrapping as an objection to Reliabilism. Vogel draws the conclusion that knowledge requires justification—of an evidentialist sort. But in fact, the bootstrapping problem generalizes to evidentialist theories as well.[21] That is, it generalizes to any evidentialist theory that allows for basic knowledge.

Consider again evidentialist foundationalism: According to that view, I can know the table is red on the basis of its looking red, even though I have no prior evidence that something's looking red is a reliable indication that is

[20] This was suggested to me by Jim Pryor—in correspondence.

[21] This was noticed by Bergman as well.

red. But then once I know the table is red, I can appeal to that fact in reasoning. A little introspection will tell me that the table appears red. So now I know that the table looks red and that it is red. So I now have *some* evidence that something's looking red is a reliable indication that it is red. And by taking a few more looks, I can acquire more evidence.

I take it that this result is no more palatable here than it was in the case of non-evidentialist views like Reliabilism. So the evidentialist who allows for basic knowledge faces the problem of easy evidence, as well. And assuming that by amassing enough such evidence, one could know that something's looking red is a reliable indication that it is red, the evidentialist faces the problem of easy knowledge as well.

One consequence of the problem of easy evidence is that it can be difficult to get a consistent statement of a view that denies the KR principle and so allows for basic knowledge. Consider e.g., what the reliabilist wants to say: One can know that p, provided that one's belief that p is produced by a reliable process, even though one has no evidence that one's process is reliable. Presumably this is supposed to hold true even for beings with self-knowledge—beings with knowledge of their own belief states. But if the bootstrapping argument is correct, anytime one knows P and knows that one believes P, one has evidence that the process that produced one's belief in P is reliable. But then if the process that produced one's belief in P is reliable, one can't know P without having evidence that the process is reliable.

Now perhaps reliabilism is only committed to holding that one can know on the basis of a reliable process without having any *prior* evidence that the process is reliable. But then the reliabilist would have to hold that once you come to know P, simply reflecting on the fact that you believe P creates evidence that the process that produced P is reliable. And similar considerations hold for evidentialist foundationalism.

IV Bootstrapping and the No Self-Support Principle

We have seen that any view that allows basic knowledge (along with a modest amount of self-knowledge) will also allow for bootstrapping. When Vogel points out that Reliabilism allows for bootstrapping, he says that he "assumes bootstrapping is illegitimate", and so Reliabilism must be rejected. Vogel is here making an intuitive appeal to which I, myself, am most sympathetic. I assume most others are as well.

Fumerton takes the illegitimacy of bootstrapping to illustrate a more general point.

> If we understand epistemic concepts as the externalists suggest we do, then there would be no objection in principle to using perception to justify reliance on perception, memory to rely on memory, and induction to justify reliance on induction. But there is no philosophically interesting concept of justification or knowledge that would allow us to use a kind of reasoning to justify the legitimacy of using that reasoning.[22]

Since Fumerton seems to be talking about perception and memory, as well as reasoning, Michael Bergmann takes Fumerton to be endorsing what Bergmann calls "*The No Self-Support Principle*":

> NSS: One cannot obtain a justified or warranted belief that a belief source S is trustworthy by relying even in part on source S.[23]

Bergmann interprets Fumerton and other internalists as arguing that externalist theories violate NSS, and because of this, externalist epistemic properties fail to satisfy the *Skeptical Controversy Condition*:

> SCC : being an epistemic property whose exemplification is at issue in the controversy between the skeptics and the non-skeptics.[24]

According to Bergmann, the internalist then argues that any property that fails to satisfy SCC is philosophically uninteresting—either it's false or changes the subject. Bergman argues in defense of externalism that if the internalist argument is sound, then no epistemological theory, internalist or externalist, satisfies SCC. So there are no grounds for preferring internalism to externalism

But it is not at all clear that the problem with views that allow bootstrapping is that they violate NSS and so fail to satisfy SCC. It's not clear to me how to decide if a view satisfies SCC. Moreover it is not clear that the problem with views that permit bootstrapping is that they violate NSS. Bootstrapping *is* a way of coming to know a belief source is reliable by in part relying on that very source. And as Vogel notes, there is a very strong intuition that bootstrapping is unacceptable. Fumerton takes this intuition to be an instance of a more general intuition expressed by NSS. But it does not follow from the illegitimacy of bootstrapping, that *any* way of coming to know a belief source is reliable by in part relying on that very source is objectionable. That would follow only if any view that violates NSS also sanctions bootstrapping. We have seen that theories which deny NSS by allowing for basic knowledge also allow for bootstrapping. But I shall argue that there are ways of denying NSS that do not allow for bootstrapping. It remains to be seen whether these alternatives which violate NSS without endorsing bootstrapping are plausible.

---

22 Fumerton, p. 180

23 Bergmann, p. 168. Bergmann couches the statement more generally, in terms of epistemic properties.

24 Bergmann, p. 167.

V *A Priori* Knowledge of Reliability

We have seen two related ways in which the problem of easy knowledge arises for theories that allow for basic knowledge—by the closure principle, and by bootstrapping. So to avoid these problems, we must accept the KR principle and eschew basic knowledge. Is there a plausible view of this kind? And must any such view accept NSS? Must such a view hold that knowledge sources can play no role in legitimizing themselves?

One strategy is to hold that we know *a priori* that our belief sources are reliable. This would allow for the possibility of knowing e.g., that perception is reliable prior to gaining knowledge by means of perception, thereby avoiding basic perceptual knowledge. Similarly one could also know that memory and induction are reliable without relying on those faculties, thereby avoiding basic memory and basic inductive knowledge. But in order to avoid *any* basic knowledge, such a view would have to reject the NSS principle. Because to avoid basic *a priori* knowledge, the view must hold that we have no *a priori* knowledge prior to knowing that *a priori* belief formation is reliable. But if we cannot have empirical knowledge prior to knowing, *a priori*, that perception is reliable, then surely we will have to know *a priori* that *a priori* beliefs are reliably acquired. And this violates NSS.

One might object to the claim that one can know *a priori* that *a priori* belief formation is reliable. But I want to consider the question of whether it plausible to hold that we know *a priori* that any cognitive processes is reliable. The reliability of our cognitive processes is a contingent matter. And so to know *a priori* that they are reliable would be to have contingent *a priori* knowledge. Surely this should give us pause.[25]

But recently, John Hawthorne has argued that we do have contingent *a priori* knowledge relevantly like this.[26] Suppose a statement of the following form is a true statement: It is rational to believe P on the basis of E, or E constitutes sufficient grounds to believe P. And let us suppose, that we can come to know such statements *a priori* (as we must if there is any point in doing traditional epistemology). If so, Hawthorne notes, it would be *a priori* rational to believe the contingent statement *If E then P*. But then such a statement seems to be a good candidate for contingent *a priori* knowledge. So on the supposition that there are some knowable truths about rational grounds for belief, it looks not so implausible to hold that there is contingent *a priori* knowledge of this kind.[27]

[25] Note that this case does not fit the reference-fixing model of the contingent apriori endorsed by Kripke and others.

[26] Hawthorne.

[27] Why the hesitation? Because the details about causal etiology of *a priori* belief are obscure enough that it's unclear whether the *a priori* belief can be related to its object in the right way to produce knowledge.

Now consider evidentialist foundationalism Such a view tells us that X's looking red constitutes sufficient grounds for believing X is red. But then by Hawthorne's argument, perhaps, for any X, I can be *a priori* (*prima facie*) warranted in believing (and if things work out, *a priori* know) that if X looks red, then X is red.[28] Indeed this might be the most plausible way to understand evidentialist foundationalism. Thus on the basis of this *a priori* knowledge, when something looks red to me, I could thereby know, or at least be warranted in believing, that my color vision was operating accurately.

Would such a view avoid basic knowledge? No particular exercise of my color vision would yield knowledge prior to my knowing that this particular exercise was accurate. But knowing that a particular exercise of my color vision is accurate is not to know that my color vision is reliable. So on this view, I would have knowledge via color vision prior to knowing that color vision is reliable. So I would have basic beliefs.

Is the view subject to the problems of basic beliefs? The view does allow me to use color vision to gain knowledge that color vision is reliable. As we have seen, with each exercise of my color vision, I gain knowledge that my color vision worked accurately on that particular occasion. Thus by repeated use of my color vision I can accumulate evidence that color vision is reliable and thereby come to know color vision is reliable.

But so far, this is just a further violation of NSS. Is it an instance of bootstrapping? This is hard to assess since we do not really have a precise definition of bootstrapping. All we have is an intuition that certain ways that some theories violate NSS are intuitively objectionable. The way the *a priori* view violates NSS is structurally different from the way Reliabilism and (non *a priori*) evidentialist foundationalism allow such violations—what we have called "bootstrapping". Because on the *a priori* view, for any X, I have prior (*a priori*) warrant for believing that if X looks red, then X is red. Does this make it unobjectionable to then infer from particular exercises of my color vision, that color vision is reliable? I do not have a clear intuition here.

Does the *a priori* face the closure problem? Suppose I know *a priori* that if X looks red, then X is red. I will know this prior to having any evidence against the possibility that if it looks red, then it is white with red lights shining on it. But this will not prevent me from knowing that if X looks red, it is not white with red lights shining on it. Indeed closure secures this knowledge for me—*a priori*. But it seems implausible to say that I could know *that, a priori*. Again, this would make my inductive evidence against the likelihood of there being red lights shining on X, irrelevant to my know-

[28] I'm not *a priori* warranted in believing that for any X, if X looks red then it's red. I know that's false. Rather, for any X, I'm *a priori* warranted in believing that if X looks red, then X is red.

ing that if X looks red, X is not white with red lights shining on it. Suppose my son worries that when X looks red, it's really white with red lights shining on it, I could reply, that if X looks red, then it is red (after all I know this—*a priori*), and thus if X looks red, X is not white with red lights shining on it. The absurdity of such reasoning indicates that I did not really know that if X looks red, then it is red, in the first place.

None of this in itself counts against Hawthorne's argument for contingent *a priori* knowledge deriving from our knowledge of epistemic principles. Rather, it at most shows that if Hawthorne's argument is correct, then epistemic principles of the sort that license basic knowledge, e.g., the foundationalist principle that *X's looking red constitutes sufficient grounds for believing X is red,* are not true (or at least not knowable) epistemic principles.

VI Holism

The second option for avoiding basic knowledge is to hold that epistemic support is to a certain extent, holistic.[29] According to this view, in the initial stages of cognitive development, our perceptual beliefs do not count as knowledge, nor does any belief we may have regarding the reliability of our faculties. Gradually, as we acquire more and more sensory evidence, thereby accumulating a relatively large and coherent set of beliefs, those beliefs, including the belief that our cognitive faculties (perception, memory, reasoning) are reliable become knowledge.

So our initial sensory evidence is not by itself sufficient for us to know things about the world.[30] Only after such evidence enables us to achieve a sufficiently coherent picture of the world and our relation to it do we acquire such knowledge. But neither our first-order knowledge of the world gained by our cognitive faculties, nor our knowledge of the reliability of our faculties is prior to the other. Rather our first-order beliefs and our beliefs regarding the reliability of our faculties stand in a mutually supporting relationship. But, unlike a pure coherence theory, the mutual support relations among beliefs are not by themselves sufficient for those beliefs to be knowledge. In order to be knowledge these beliefs in the coherent set must be supported by sensory evidence.[31] The status of the beliefs as knowledge depends both on their origins in sensory experience and in their mutual support relations.

It will be useful to compare the holistic support view to evidential foundationalism. Both views hold, e.g., that our knowledge of the reliability of perception depends on sensory evidence. Each view thus rejects NSS. According to foundationalism, perceptual evidence yields an initial fund of perceptual knowledge on the basis of which we eventually learn enough about the world and our relation to it that we can come to know that perception is reliable. So our perceptual evidence delivers knowledge—basic knowledge—of the world, and this knowledge of the world delivers knowledge of the reliability of perception. The holistic support view differs only in the way it construes the epistemic priority relations. While our perceptual evidence produces perceptual beliefs, and such beliefs produce, or at least psychologically sustain, our beliefs about the reliability of perception, the perceptual beliefs do not become knowledge prior to the belief regarding the reliability of perception becoming knowledge. So there is no asymmetric epistemic support relation between our perceptual beliefs and our belief in the reliability of perception.

Does holism avoid basic knowledge and the attendant problem of easy knowledge? Although holism does not allow us to acquire e.g., perceptual or memory knowledge prior to our acquiring knowledge of the reliability of perception and memory, it does not follow that there are no basic beliefs. Because on the holistic support view, the source of one's perceptual knowledge, memory knowledge, and indeed all one's initial empirical knowledge is holistic support—which comprises perceptual evidence along with coherence relations among beliefs, including the belief that one's knowledge sources are reliable.[32] This raises the question of whether this latter belief that one's knowledge sources are reliable applies to holistic support itself. If it does not, then all my knowledge, perceptual, memory, etc., will be basic, in my sense. I will have attained it prior to my knowledge that its source, viz., holistic support, is reliable. And there would be nothing to prevent one from bootstrapping this basic knowledge into knowledge of the reliability of holistic support.

To avoid this problem, the proponent of holistic support must argue that we can know on the basis of holistic support that holistic support is reliable.[33] Moreover such knowledge cannot be subsequent to any other knowledge acquired by means of holistic support. In order to know that holistic support is reliable, I would have to know I am not systematically deceived, e.g., I would have to know such things as that I am not a brain-in-a-vat. Notoriously, it is very difficult to give some account of how one might know that such global deceptions are not occurring. Perhaps holistic support will have to involve some pragmatic criterion like simplicity.

---

29 This view has been advocated in various forms by various philosophers—Harman, Lehrer, Sellars, Sosa (1997) (forthcoming), Vogel (1970).

30 If sensory evidence is default belief, then the view would be that these beliefs and the fact that we have them supports their truth.

31 Harman, Vogel, Sosa (1997).

32 Of course I don't mean that holistic support is the psychological source of one's beliefs. Rather I mean that holistic support is what makes the beliefs knowledge, and in that sense, is the source of one's knowledge.

33 Of course, typically, we don't have beliefs regarding holistic support. But so long as we have sufficient reason to believe holistic support is reliable and that reason plays a role in sustaining other beliefs, then we can avoid basic beliefs.

Of course specifying what these criteria are and how they work has proven to be elusive.[34] But this is not a burden unique to the holistic support view. If views that allow for basic knowledge are to explain our knowledge of the world and how we are cognitively situated in it, they must provide some account of how we know global deceptions are not occurring. And allowing for basic knowledge does not, by itself, provide for such an account. Suppose we allow that we have basic perceptual and memory knowledge. How do we get from that to the conclusion that we are not systematically deceived e.g., that we are not brains-in-a-vat? We have already discussed two illegitimate ways to do so—by noting that one's basic knowledge entails that one is not systematically deceived, and by bootstrapping. But even if we were to countenance these ways of knowing, they would only give us knowledge that we are not *deceived* brains-in-a-vat. A fact P about the world that I might know by perception does not entail that I am not a brain-in-a-vat, only that if I am, then P. And the inductive evidence I gain by bootstrapping only establishes again that when I believe P, P is the case. But in itself, this does not rule out that I am not the "victim" of a benevolent demon (or vatmeister)—i.e., that I am not the victim of systematic veridical hallucinations. So I can at best know that if I am subject to systematic hallucinations, my beliefs are still true.

Do BKS views have to provide an account of how we know we are not subject to systematic hallucinations—veridical or otherwise? Such views are supposed to explain how we acquire knowledge of the reliability of our faculties. So if such views can explain how I know that if I am subject to systematic hallucinations, then my beliefs are still true, do I thereby know my faculties are reliable, in a way that would secure knowledge? Since I myself do not have reliabilist intuitions, I do not know how a reliabilist would answer this question. Certainly by attaining such knowledge, one will not have attained "a broad view of oneself and one's environing world" (Sosa) which BKS theories are supposed to deliver. To attain such a view, I would have to know I am not the victim of systematic veridical hallucinations. And this means that BKS theories are confronted with the same kind of difficulty faced by holistic support theories—ruling out the possibility of radical skeptical scenarios. So if holistic support theories must rely on pragmatic criteria, so presumably must BKS theories.

Moreover even if BKS theories could somehow deliver knowledge that I am not a brain-in-a-vat, in a way that holistic support theories could not, there remain other skeptical alternatives. Consider the possibility that I am deceived by a demon who after I observe an object destroys and creates a duplicate, or the possibility that all apparent causal relations in the world are mediated by the intervention of a demon. The mere fact that I have basic

[34] Perhaps the most successful attempt to specify occurs in Vogel (1990).

perceptual knowledge does not by itself enable me to know that these scenarios do not obtain. And in order for me to know that I am reliable, again, to gain a broad view of myself and my environing world, I must know that these possibilities do not obtain. Again this means that BKS theories will presumably have to appeal to pragmatic criteria. So the problem of how we know we are not victims of radical skeptical deceptions does not provide any basis for preferring BKS theories to holistic support theories.

On the supposition that holistic support theories can account for how we know we are not systematically deceived, we can hold that we know by holistic support, that holistic support is reliable. Thus, the knowledge we gain by holistic support will not be prior to our knowledge that holistic support is reliable. Thus holistic support will not allow for basic knowledge and so can avoid the problem of easy knowledge.

VII Two Problems: Basing and Unreflective Subjects

I have just provided a sketch of a holistic support view. But let me mention two worries about the prospects for developing a workable developed version of the view.

The first problem concerns the basing relation. There is a well-known distinction between merely having good reasons for believing P and believing P for good reasons. This latter notion requires that one's good reasons for believing be psychologically related in the right way to one's believing P, i.e., that one's believing P be *based* on one's good reasons for believing P.

On the holistic support view a significant chunk of one's beliefs become knowledge in virtue of mutual support relations, e.g., a chunk of one's beliefs about the world and one's belief that perception is reliable. But this entails that each member of the mutually supporting set of beliefs must be based on the other members of the set. Many epistemologists think that the basing relation is at least in part a causal notion, in some extended sense of the term (comprising causal sustenance, causal overdetermination, pseudo-overdetermination, etc.) If so, then the holistic support view must hold that the members of the mutually supporting set of beliefs are in some epistemologically important way causally related to one another.[35]

How plausible is it to suppose that they are related in this way? I don't know. The answer to this question will depend on the precise nature of the basing relation and the facts of psychology. But clearly there is cause for concern here.

The second problem pertains to the holistic view's contention that we do not have any perceptual knowledge until we know that perception is reliable. But knowing one's perceptual faculties are reliable is a relatively sophisticated cognitive task—one that would seem to be beyond the reach of unreflec-

[35] Pollock (1979).

tive beings like small children and animals. Yet it is hard to deny that such beings have some perceptual knowledge. For example, a 2 year-old certainly can know that there is something red before him. Yet presumably s/he does not know that his/her perceptual faculties are reliable. The apparent fact that unreflective beings have knowledge suggests that we must allow for basic knowledge.

If I'm right about the problems confronting BKS theories, then there would seem to be a dilemma. On the one hand, if we allow basic knowledge, we face the problem of easy knowledge. On the other hand, if we give up on basic knowledge, then we end up denying that unreflective beings like small children have simple perceptual knowledge. Of course it is difficult to say exactly what unreflective beings like 2 year olds know. What exactly is the content of their beliefs? But it does seem hard to resist the claim that there is something that they know.

In recent work, Ernest Sosa, following his reading of Descartes, distinguishes between unreflective (animal) knowledge and reflective knowledge.

> ...we can more generally distinguish between *animal* knowledge which requires only that one track reality, on one hand, and *reflective* knowledge, on the other, which in addition requires awareness of how one knows, in a way that precludes the unreliability of one's faculties.[36]

On Sosa's view, knowledge has a basic knowledge structure with animal knowledge playing the role of basic knowledge.[37] Because of this, his view is subject to the problem of easy knowledge. This is the dilemma with which we began this section. All the same, the idea of animal knowledge is attractive as a way to account for our intuition that unreflective being know things. But perhaps we could allow for the existence of animal knowledge and avoid the pitfalls of easy knowledge.

Non-reflective beings, by definition, do not form beliefs about how things appear to them, or indeed any beliefs about their cognitive faculties. And it follows from this fact alone that non-reflective beings do not have knowledge about the reliability of their cognitive faculties. But it does not follow that we can allow for basic knowledge in unreflective beings and avoid the problem of easy knowledge. Some non-reflective beings, eventually, if all goes well, become reflective beings. And there would be nothing to stop them at that point from bootstrapping their animal knowledge into knowledge of the reliability of their faculties.

Let me tentatively suggest another strategy. Animal knowledge represents a relatively minimal cognitive achievement—something that relatively unsophisticated beings routinely obtain. Perhaps we should take seriously the

[36] Sosa (forthcoming).

[37] On Sosa's view, once you achieve knowledge of the reliability of your faculties, your perceptual knowledge shifts from being animal knowledge to being reflective knowledge.

idea that animal knowledge is a different kind of knowledge. In light of this, we could view it as differing from full-blown reflective knowledge with respect to its role in reasoning. For one thing, we could suppose that animal knowledge is not closed under known entailment. Indeed, if animal knowledge requires, as Sosa supposes, nothing more than tracking, it will follow that animal knowledge does not respect closure. Thus it will not follow from the fact that one has animal knowledge that X is red, that one has animal knowledge that X is not white and illuminated by red lights. This would not do violence to our strong intuition supporting closure, since that intuition can still be taken as revealing something deep about the nature of reflective knowledge. So if animal knowledge does not obey closure, then we can avoid the problem of easy knowledge engendered by the closure principle.

To prevent bootstrapping to easy knowledge, we would need to suppose that instances of animal knowledge cannot combine individually (non-holistically) with self-knowledge to generate inferences. So I can not combine my animal knowledge that there is something red before me with my knowledge that my belief is produced by color vision in order to infer that my color vision is working correctly.

This would still allow that animal knowledge could combine with other animal knowledge to support rudimentary inductive inferences. Just as we are tempted to say that some unreflective beings have perceptual knowledge, we are also tempted to say that they can know by induction. For example when the doorbell rings, the two year old can know there is someone at the door.

Is this an *ad hoc* view? Perhaps it is—to some degree. But if my arguments are plausible, then there are opposing demands on our epistemological theorizing. On the one hand, the naturalness with which we attribute knowledge to non-reflective beings, as well as (for many) the intuitive attractiveness of BKS theories in general, shows that there is something very plausible about the idea that knowledge can arise from conditions significantly more meager than those required by holism. But the problem of easy knowledge shows that any such conception will lead to unacceptable consequences regarding the acquisition of knowledge about the workings of our cognitive faculties. We can accommodate these seemingly conflicting demands by supposing, as does Sosa, that there are two kinds of knowledge, one relatively primitive, the other more complex. Given that animal knowledge is a more primitive form of knowledge, it is not unreasonable to suppose that it lacks some of the inferential versatility of the more complex, reflective knowledge. And if we suppose it to be limited in just the ways I have suggested, we can account for our intuition that unreflective beings have knowledge and still avoid the problem of easy knowledge.[38]

[38] I would like to thank Richard Feldman, John Hawthorne, Jim Pryor, and Jonathan Vogel for helpful discussion. In a number of papers, Vogel has raised issues that anticipate, and

References

Alston, William, *Epistemic Justification* (Cornell University Press, Ithaca, 1989)

Bergmann, Michael, "Externalism and Skepticism", *The Philosophical Review,* 2000

Bonjour, Laurence, *The Structure of Empirical Knowledge* (Harvard University Press, Cambridge, MA, 1985)

Cohen, Stewart, "How to be a Fallibilist", *Philosophical Perspectives* 2, James Tomberlin (ed.) (1988)

Cohen, Stewart, "Contextualism, Skepticism, and the Structure of Reason", *Philosophical Perspectives* 13, James Tomberlin (ed.) (1999)

Dretske, Fred, "The Pragmatic Dimension of Knowledge", *Philosophical Studies*, 40, 1981

Fumerton, *Metaepistemology and Skepticism* (Rowman & Littlefield, Lanham, Maryland, 1995)

Goldman, Alvin, *Epistemology and Cognition* (Harvard University Press, Cambridge, MA, 1986)

Harman, Gilbert, *Thought (*Princeton University Press, Princeton, 1974)

Hawthorne, John, "Deeply Contingent *A priori* Knowledge", *Philosophy and Phenomenological Research*, 65, 2002

Klein, Peter, "Skepticism and Closure: Why the evil genius argument fails", *Philosophical Topics,* 23, 1995

Lehrer, Keith, *Knowledge* (Oxford, 1974)

Nozick, Robert, *Philosophical Explanations* (Harvard University Press, 1981)

Pollock, John, "A Plethora of Epistemological Theories", *Justification and Knowledge,* George Pappas (ed.) (Reidel, Dordrecht, 1979)

Pollock, John and Cruz, Joseph, *Contemporary Theories of Knowledge* (Rowman and Littlefield, Totowa, NJ, 1986)

Pryor, James, "The Skeptic and the Dogmatist", *Noûs* 34, 2000

Sellars, Wilfred, "Givenness and Explanatory Coherence", *Journal of Philosophy*, 70, 1973

Sosa, Ernest, "Reflective Knowledge in the Best Circles", *Journal of Philosophy*, 94, 1997

Sosa, Ernest, "Virtue Epistemology" in *Epistemology: Internalism Versus Externalism* (co-authored with Laurence Bonjour) forthcoming

van Cleve, James, "Foundationalism, Epistemic Principles, and the Cartesian Circle", *Philosophical Review*, 1979

van Cleve, James, "Reliability, Justification, and the Problem of Induction", *Midwest Studies in Philosophy* (French et al., eds.) 1984

are closely related to, what I am calling "the problem of easy knowledge". See Vogel (1993), (1999), (2000).

Vogel, Jonathan, "Cartesian Skepticism and Inference to the Best Explanation", *Journal of Philosophy,* 1990

Vogel, Jonathan, "Reliabilism Leveled", *Journal of Philosophy*, 2000