

Chapter 4

Is Knowledge Easy – or Impossible? Externalism as the Only Alternative to Skepticism

James Van Cleve

In a number of articles I have defended the idea that we may obtain knowledge of the reliability of our cognitive faculties by using those very faculties – a procedure that to many looks objectionably circular (Van Cleve 1979, 1984, 1999). An important part of my strategy has been to invoke a kind of epistemic externalism, which implies (or so I argue) that there is nothing objectionably circular about the contested procedure after all. To this the reaction of some has been: so much the worse for externalism. Any theory that implies that we can have knowledge of the reliability of a faculty by using that very faculty makes knowledge too easy.¹ In what follows I press the suggestion that the only alternative to such externalism may be skepticism. If correct, this is a significant result, for those who object to the circularity sanctioned by externalism do not generally wish to embrace skepticism. More typically, they believe that there is some third alternative to the easy knowledge of the externalist and the unattainable knowledge of the skeptic. But the existence of any such third way is precisely what I shall question.

Varieties of Externalism

I begin with a rough taxonomy of various of the internalisms and externalisms in epistemology and an indication of which variety of externalism is crucial for my project.² I characterize the internalisms positively and define each externalism simply as the denial of the corresponding internalism. 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. Theories that fit under this slogan may vary in two dimensions – according to the *object* of the higher-order knowledge and according to the *nature* of the higher-order knowledge. As regards the object of the higher-order knowledge, two main possibilities are (1) that the knowledge-making or knowledge-conferring factor *obtains*, and (2) that the factor, whatever it may be, has ‘justificatory efficacy’ (to use a phrase of Alston’s (1989b: 227–49)) – for example, that it does indeed confer knowledge, or that it is a reliable source of true beliefs.³ As regards the nature of the higher-order knowledge, we get four possibilities, depending on how we

answer two questions. First, must the higher-order knowledge be (a) actual, or may it be (b) merely potential, that is, a matter of the subject's being in a position to know? Second, must the knowledge arise or be obtainable (i) in some prescribed way, typically specified as by reflection alone, or may it arise (ii) in any way? These distinctions are enough to give us eight internalist theses. We get further possibilities by substituting 'justification' or 'justified belief' for 'knowledge' in various places above, but I shall leave that to the reader.⁴

The externalism that will figure importantly in what follows is an externalism that denies that a subject need actually know, by reflection or otherwise, that the factors justifying his belief or giving him knowledge are reliable indicators of truth. In other words, it affirms there are knowledge-making factors that do their work of giving subjects knowledge even if the subjects do not know anything about their reliability. In terms of the distinctions above, it is a denial of internalisms 2ai and 2aii, at least if 'justificatory efficacy' is held to involve reliability. Externalism of this variety is quite compatible with internalism of the 2bii variety, which merely affirms that the justificatory efficacy of knowledge-making factors must be *knowable*. Indeed, the 'easy knowledge' objection is precisely the objection that externalism makes such efficacy too easily knowable.

Reliabilism is one clear example of a theory that is externalist in my indicated sense. For a reliabilist, the fact that a subject's belief has been produced by a reliable mechanism suffices for the belief's being justified (and in favourable cases, for its being a piece of knowledge) whether the subject has knowledge of this fact or not; there is no requirement of higher-order knowledge.

But reliabilism is not the only variety of externalism. It is worth pointing out that some non-reliabilist views, including views that could be considered paradigms of internalism, are nonetheless externalist in the sense that will be crucial here. Here is Chisholm: 'the concept of epistemic justification is ... *internal* ... in that one can find out directly, by *reflection*, what one is justified in believing at any time' (1977: 7).⁵ Given that what one is justified in believing is a function of two things, certain justifying states and the relations whereby they confer justification on what they justify, this makes Chisholm an internalist in two of the senses distinguished above, 1bi and 2bi: if one is justified by some factor *f* in believing *p*, one must be able to know by reflection both that *f* obtains and that *f* does confer justification on belief in *p*. Now this arguably implies that justifying factors do their work even if they are not reliably connected with the truth of what they justify, since reliability relations do not seem to be the sort of thing that can be known to obtain by reflection alone.⁶ So for Chisholm, justifying factors need not be reliably connected with what they justify.⁷ A fortiori, justifying factors need not be *known* to be reliably connected with what they justify. And that means that Chisholm is an externalist in the sense that matters here: there are sources of justification or knowledge that deliver their goods even if the subject does not know they are reliable.

Easy Knowledge

Externalism of the sort just indicated makes it possible for subjects to come to know

that a faculty is reliable by reasoning from premises delivered by that very faculty. For example, reliabilist externalism makes it possible to argue for the reliability of sense perception using premises derived from sense perception. Such an argument might take the following form, which Alston (1986b: 9) has dubbed a 'track-record argument':

- 1 At *t*1, I formed the perceptual belief that *p*, and *p*.
- 2 At *t*2, I formed the perceptual belief that *q*, and *q*.
(and so on)
- C Therefore, sense perception is a reliable source of belief.

In such an argument we consider a wide sample of perceptual beliefs, note that the vast majority of them have been true, and conclude by induction that perception is reliable.⁸

There is a strong tendency to regard such arguments as circular. The reason for this, as is plain, is that the second conjunct in each premise could be known only on the basis of sense perception. It would not have to be known on the basis of the very exercise of perception that initially prompted the belief – instead it could be known through subsequent corroborative exercises of perception, or even exercises of another faculty (for example, one's natural credulity in believing a witness). But it remains the case that to ascertain the truth of the premises, one would have to make use of one's general perceptive faculties (for example, in coming to know that the witness uttered certain words). So in running through a track-record argument, one is indisputably using the deliverances of a faculty to argue for the reliability of that very faculty.

I maintain, however, that on reliabilist principles (and also on the more broadly externalist principles identified above), there is nothing objectionably circular about such arguments. This is so for two reasons jointly. First, an argument is vitiatingly circular – that is, circular in a way that makes it inept as a means for acquiring knowledge of its conclusion – *only if knowledge of the conclusion is a precondition of acquiring knowledge of its premises.*⁹ Second, if reliabilism is correct, knowledge of the reliability of sense perception is *not* a precondition of using sense perception to gain knowledge, so *knowledge of the track-record argument conclusion is not a precondition of acquiring knowledge of its premises.* The senses need merely be reliable in order to be means of acquiring knowledge, whether a subject knows them to be or not. If they are reliable, then a subject may acquire knowledge of various environmental facts, including the *ps* and *qs* in the argument above, without any antecedent knowledge of the conclusion. So the argument is not vitiatingly circular. On reliabilist principles, it is not subject to any reproach.

And yet it seems to many that there is something illegitimate about such arguments nonetheless. Jonathan Vogel has made it an objection to reliabilism precisely that it does permit such arguments. He parodies them with the example of Roxanne, who drives a car with a highly reliable gas gauge and who

believes implicitly what her gas gauge says, without knowing that the gauge is reliable. But she undertakes the following, admittedly curious, procedure. She looks at the gauge often. Not only does she form a belief about how much gas is in the tank, but she also

takes note of the state of the gauge itself. So, for example, when the gauge reads 'F', she believes that, on this occasion, the tank is full. She also believes that on this occasion, the gauge reads 'F'. Moreover, Roxanne combines these beliefs; she believes

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

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

(Vogel 2000: 613)

The reader can guess how the example goes on from here. Roxanne repeats this procedure over and over, forming many beliefs of the form

On this occasion, the gauge reads 'X' and X.

Putting all these bits of information together, she eventually infers by induction 'The gauge is reliable', and from that, together with her knowledge of how she is arriving at her own beliefs, she infers in turn that the process by which her beliefs about tank levels are formed is reliable. She thus promotes beliefs that were formed by a reliable process into *knowledge* that those beliefs were formed by a reliable process.

Vogel takes it for granted that Roxanne's procedure, which he calls bootstrapping, is an illegitimate way of gaining knowledge of reliability. But he also observes that reliabilism apparently ratifies bootstrapping. He concludes that reliabilism should be rejected. He goes on to suggest that we can avoid having to endorse bootstrapping if we adhere to the traditional view (now often called evidentialism) that reasons or justification are necessary for knowledge. If justification is necessary for knowledge, he says, Roxanne does not have knowledge of the premises involved in the bootstrapping argument:

If Roxanne has no independent reason to believe that the position of the needle on the gauge is reliably correlated with how much gas is in the tank, I cannot see that she is justified in believing that her tank is full when she looks at the gauge. So the traditionalist can say that Roxanne does not know (20) or (21) ['On this occasion, the gauge is reading accurately'], since she lacks justification for both. There are no bootstraps for her to pull on.

(Vogel 2000: 622)

But evidentialism as such is not enough to thwart bootstrapping. An evidentialist few will also permit bootstrapping, just so long as it says there are sources that confer knowledge or justification without any requirement that the reliability of the source be known. As I observed above, Chisholm's evidentialism is one such view. In fact, in order for the bootstrapper to get off the ground, all we need allow is that perception is a source of knowledge independent of knowledge of its own reliability, and similarly for introspection, memory and induction. For by introspection one may know that perception is the source of one's belief in *p*; by perception itself

(though not necessarily the very same exercise of it) one may know *p*; by combining these two, one may know that perception was truthful in a given instance; by memory, one may then know that perception was truthful in a great many instances (and deceptive, we may assume, in few or no instances); and by induction, one may then conclude that perception is reliable. It seems that one could take issue with this result only by questioning the sufficiency of perception, introspection, memory or induction to give knowledge – or at least their sufficiency in the absence of any knowledge of their reliability.

The point that evidentialism no less than reliabilism gives rise to a kind of 'easy knowledge' has been very well made by Stewart Cohen (2002). Cohen discusses two varieties of 'easy knowledge', one variety obtained through the use of epistemic closure principles (which I do not consider here) and the other obtained through track-record arguments, such as I have discussed. He argues that any epistemology will give rise to 'easy knowledge' just so long as it denies the following principle:

(KR) A potential knowledge source *K* can yield knowledge for a subject *S* only if *S* knows *K* is reliable.

I have myself argued in an earlier article for a similar conclusion regarding induction – that the use of induction to establish the reliability of induction (clearly a form of bootstrapping) is legitimate not only for reliabilism, but also for what I called 'apriorist' or traditional epistemology. Allow me to repeat what I said at the end of that article:

What makes the reliabilist solution [to the problem of induction] work is its tenet that reliability is *sufficient* for justification; what makes the apriorist solution work is its tenet that reliability is *not necessary* for justification. Hence, in order to reject both solutions, one would have to maintain that reliability is *necessary but not sufficient* for justification. And in that case the following question would become pressing: what is the *x* that must be added to reliability to obtain justification?

(Van Cleve 1984: 565)

The answer I thought would immediately suggest itself is this: what must be added is *knowledge* of reliability. With that answer, we arrive, in effect, at *KR*: a source *K* yields knowledge for a subject *S* only if *S* knows that *K* is reliable.

Externalism says just the opposite: that a source can deliver knowledge to a subject even if the subject does not know the source is reliable. This would be true for a reliabilist for whom the reliability of a source suffices by itself to make its deliverances knowledge; it would also be true for an evidentialist like Chisholm or Descartes, who holds that there are sources (for example perception and introspection) that make their deliverances *prima facie* justified (and thus in favourable circumstances knowledge) by virtue of necessary principles of justification, no matter what reason the subject has to believe the sources reliable.

It is externalism's denial of *KR* that makes bootstrapping possible. If a source can deliver knowledge to one who does not yet know the source to be reliable, then the subject can use the deliverances of that source to amass evidence from which he can infer by induction that the source is indeed reliable. Hence it is that some will see the affirmation of *KR* as required to prevent bootstrapping.

KR as Leading to Skepticism

And yet the affirmation of *KR* threatens to land us in skepticism.¹⁰ At least it would do so if we make three assumptions. First, we assume that what *KR* lays down as a necessary condition of knowledge through *K* is a *prior* condition of it, that is, that *S* can know through *K* that *p* only if *S* first knows that *K* is reliable. Second, we assume that knowledge of the reliability of a source can come about in one way only, namely, by inference from premises obtained from that very source. In other words, we must know that various of the particular deliverances of a source are true before we can know that the source is reliable; and if the source is an ultimate source, like perception or memory, there will be nothing but the source itself to vouch for its own deliverances.¹¹ Putting these points together, we find ourselves in the following predicament:

- 1 We can know that a deliverance of *K* is true only if we first know that *K* is reliable.
- 2 We can know that *K* is reliable only if we first know, concerning certain of its deliverances, that they are true.

If (1) and (2) are both true, skepticism is the inevitable consequence. Clearly, if we cannot know either of two things without knowing the other *first*, we must remain ignorant of both of them.¹² So assuming that the term 'knowledge' has been used univocally throughout (1) and (2) – and that is our *third* assumption – skepticism prevails. There can be no knowledge through any source *K*.

Such is the *prima facie* case for my main thesis: either externalism (of the sort tantamount to the denial of *KR*) or skepticism. The case rests, however, on the three assumptions just noted, and that opens up three possible avenues of escape. Perhaps we can find a non-skeptical alternative to externalism by denying one of the assumptions; that is what I explore in the next three sections.

The Reidian Alternative

In the first alternative I wish to consider, we deny assumption (2): we deny that knowledge of the reliability of a source must be derived by inference from the particular deliverances of that source. But if not derived from knowledge of its own particular deliverances, from what other knowledge *could* knowledge of *K*'s reliability be derived? One answer is: from *no* other knowledge. Knowledge of the reliability of our faculties is epistemically basic.

One philosopher who appears to have taken this line is Thomas Reid. I list here principles 1, 3, 5 and 7 from his enumeration of the 'First Principles of Contingent Truths':

First, then, I hold, as a first principle, the existence of every thing of which I am conscious.

Another first principle I take to be, that those things did really happen which I distinctly remember.

Another first principle is, that those things do really exist which we distinctly perceive by our senses, and are what we perceive them to be.

Another first principle is, that the natural faculties, by which we distinguish truth from error, are not fallacious.

(Reid 1969: 617, 622, 625, 630)¹³

If Reid is right, we have a way of breaking the skeptical impasse. We can accept the condition that *KR* lays down for all our knowledge, but affirm that that condition is thankfully met, owing to the first principles of human knowledge.

How plausible is it to take the reliability of our various cognitive faculties as epistemically basic? 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. There are, of course, general propositions that are immediately evident (for example, all triangles have three angles), and there are also contingent propositions that are immediately evident (for example, I am now conscious). But are there any immediately evident propositions that are both general and contingent? That is a question to which I find it hard to say yes.

Let me set aside that objection to Reid, however, in order to concentrate on another worry. If *KR* is a principle governing all of our knowledge, some may wish to ask, how can there even *be* such a thing as basic knowledge? Does *KR* not tell us that any piece of knowledge depends on a further piece, namely knowledge that the first piece has a reliable source?¹⁴ And does that not imply that *no* knowledge is basic?

That line of reasoning is too hasty, however, and needs to be examined more closely. Let us suppose with Reid that it is a piece of basic knowledge that sense perception is reliable. What is the source of such knowledge? If only to give it a name, let us call it intuition. By *KR*, such intuition yields knowledge only if we know it is reliable. What is the source of *that* knowledge? This time we may answer that it is intuition itself – intuition could intuit its own reliability. So it appears that *KR* can admit basic knowledge of the reliability of a source at least in those cases in which the reliability of *K* is known through *K* itself.¹⁵

Now let's go back to the principle that *perception is reliable* – could that be basic if knowledge of it depends on the knowledge that some other faculty is reliable? Perhaps, if intuition delivers its various results with equal immediacy. Let me explain.

We should note that the fact that one proposition is justified or known only if another is does not mean that the first depends for its justification (or its being known) on the second in any sense implying that the first is not basic. Perhaps *I think* is justified only if *I exist* (or *someone thinks*) is justified; but if this is a matter of automatic transfer of justification from the first to the second, the first can still be basic. However, this does not seem to be the situation we have with the intuition that *perception is reliable* and the intuition that *intuition is reliable*: there doesn't seem to be any automatic transfer of justification from the first to the second.

What we may have, though, is a situation in which intuition confers justification

simultaneously on its primary object (that perception is reliable) and also on its own reliability. If so, we could perhaps say that our knowledge of the primary object is basic even though necessarily accompanied by knowledge of a further proposition.

The Reidian view that is emerging evidently requires that there be certain faculties or sources that deliver knowledge not only of their primary objects, but also of their own reliability. This may be what Reid himself is getting at in the following passage:

How then come we to be assured of this fundamental truth on which all others rest? [Reid is speaking of his principle 7, that our faculties are not fallacious.] Perhaps evidence, as in many other respects it resembles light, so in this also, that as light, which is the discoverer of all visible objects, discovers itself at the same time: so evidence, which is the voucher for all truth, vouches for itself at the same time.

(Reid 1969: 632)

A view something like this has also been attributed to Descartes (or suggested as an improved version of Descartes) by Mark Heller. According to Heller, Cartesian clear and distinct perception is 'an introspective mechanism that would indicate its own reliability every time it was used' (1996: 50).

In his writing on Descartes, H. A. Prichard (1950) also took up a position similar to Reid's, but more limited in the number of sources whose reliability he affirmed as basic. He said, in effect, that we can know immediately that Cartesian clear and distinct perception (which embraces introspection of our own states of mind and rational apprehension of necessary truths) is an infallible source of knowledge. One can imagine him saying: when I reflect on my current condition (of clearly and distinctly perceiving p) I see that I could not be related to any proposition as I am related now to p unless that proposition were true. This approach may be more plausible than Reid's, for perhaps it is an immediately evident necessary truth, apprehendable by a kind of intuitive induction, that a priori intuition and introspection are reliable. On the other hand, Prichard's approach, unlike Reid's, does nothing to head off skepticism about the external world, the past, or other minds.

The Coherentist Alternative

The next solution, like Reid's, accepts KR as it was stated—there is no knowledge of any proposition without knowledge of the reliability of its source. But it rejects the assumption above that what KR lays down is a *prior* condition of knowledge. Knowledge through a source K depends, indeed, on knowledge of the reliability of the source; and knowledge of the reliability of K depends in turn on knowledge of its own particular deliverances. But this is a case of mutual dependence without mutual priority; there is no need for either of the mutually dependent pieces of knowledge to be acquired *first*. Rather, they attain the status of knowledge together, once a wide enough body of supporting beliefs are all in place. This way out has been advocated by Cohen in his treatment of our problem, and it is of course the hallmark strategy of coherence theories of knowledge.¹⁶ Every item of knowledge

depends for its status as knowledge on other cognitions or beliefs that support it, but nothing is epistemically prior to anything else. A system of belief evolves into a system of knowledge as it becomes more comprehensive and coherent. As Wittgenstein said, light dawns gradually over the whole.

The coherentist alternative, unlike the Reidian alternative, denies that knowledge of reliability can be basic, for it denies that there is *any* basic knowledge. Nonetheless, there are important (and perhaps surprising) points of similarity between the two approaches. The Reidian, like the coherentist, must deny that whenever knowledge of A depends on knowledge of B , B is epistemically prior to A . (Knowledge from a source K depends for Reid on knowledge of K 's reliability, but knowledge of K 's primary objects can still be basic.) And the coherentist, like the Reidian, must admit into the fabric of knowledge a crucial element of self-support. The latter point is well brought out by Cohen, who notes that if coherence theories are to stay true to KR , they must hold that the reliability of holistic support – which in one sense is the sole source of knowledge in a coherence theory – is known through holistic support itself.¹⁷

I reserve further discussion of the coherentist alternative until we have introduced a third alternative, in some ways similar to it.

The Two-levels Alternative

We have been exploring the worry that KR , the requirement that we know something through a source K only if we know that K is reliable, leads inevitably to skepticism. It would certainly seem to do so if we render KR as

- 1 We can know that a deliverance of K is true only if we first know that K is reliable

and then add its plausible converse,

- 2 We can know that K is reliable only if we first know, concerning certain of its deliverances, that they are true.

However, (1) and (2) have skeptical consequences only if 'know' as it occurs on the left side of each has the same meaning as 'know' as it occurs on the right. The strategy for upholding KR without skeptical consequence now to be explored denies precisely that assumption. If we distinguish appropriately between two senses of 'know', or two kinds or levels of knowledge, perhaps we can uphold KR without skeptical consequences. A strategy along these lines is to be found in somewhat different forms in work by Ernest Sosa and Keith DeRose.

For some years now Sosa has been developing a distinction between animal knowledge and reflective knowledge. Animal knowledge is (roughly) reliable true belief, such as unreflective beasts may have and such as even a thermometer might have if it could be said to have beliefs at all. Reflective knowledge, on the other hand, is a meta-state that arises only once one has achieved a perspective on one's

own beliefs and their reliability. A subject has reflective knowledge of a proposition p only if, among other things, he knows that the source of his belief in p is reliable.¹⁸

So it is indeed true that we have reflective knowledge through K only if we first know that K is reliable. However, the knowledge of reliability we must have 'first' is not itself reflective; it need only be of the animal variety. As Sosa says, 'Human reflective knowledge is most likely to depend ultimately on unreflective knowledge, since we cannot climb infinite ladders of reflection' (1991: 20). Reflective knowledge of the reliability of a source depends on prior knowledge of its deliverances, and reflective knowledge of the deliverances depends on prior knowledge of reliability; but the prior knowledge is in each case itself compounded out of animal knowledge. The problem of the circle into which there is no entry dissolves.¹⁹

A structurally similar two-levels view has been advocated by DeRose (1992b) in his interpretation of Descartes. No knowledge of the reliability of clear and distinct perceptions without knowledge of the existence and veracity of God, and no knowledge of God without knowledge of the reliability of clear and distinct perceptions – such is Descartes's analogue of our problem with (1) and (2). DeRose seeks to break the circle by finding two levels of knowledge in Descartes, marked by the terms *cognitio* and *scientia*. Replying to the objection that an atheist can surely know truths of geometry, Descartes wrote:

That an atheist can know clearly that the three angles of a triangle are equal to two right angles, I do not deny, I merely affirm that, on the other hand, such knowledge on his part [*cognitionem*] cannot constitute true science [*scientia*], because no knowledge that can be rendered doubtful should be called science.

(1985: vol. 2, 39)

Descartes never elaborated this distinction in any systematic way, but as DeRose goes on to develop it, S has *scientia* of p if and only if (1) S clearly and distinctly perceives that p is true, and (2) S clearly and distinctly perceives the truth of the principle that what S clearly and distinctly perceives is true (DeRose 1992b: 224). Even that clear and distinct perception is the only source of *cognitio* for Descartes, these conditions are equivalent to the following: S has *scientia* of p iff S has *cognitio* of p plus *cognitio* of the reliability of the original *cognitio*'s source. The analogous principle relating Sosa's two types of knowledge would be this: we have reflective knowledge of p iff we have animal knowledge of p plus animal knowledge that the source of belief in p is reliable.²⁰

With either Sosa's or DeRose's distinction in hand, we may chart a way in which knowledge of reliability can arise from knowledge of particular deliverances (and vice versa) without vicious circularity. In the Cartesian setting, things might proceed like this: first, we attain various object-level *cognitiones* by letting the light of clear and distinct perception fall on various truths. From these first truths we go on to infer various other things, including most importantly the existence and veracity of God. From that in turn we infer that all clear and distinct perceptions are true. As soon as we have taken that last step, all of our *cognitiones* become pieces of *scientia*, for they are now augmented by *cognitiones* of the reliability of their source.²¹ Light dawns over the whole not gradually, but in a single flash.

For Sosa, the situation would be roughly similar. We begin with animal knowledge of various particular propositions p , q , r , along with animal knowledge that source K delivered p and so on. That gives us the data for what Russell called an 'animal induction' to the conclusion that K is a reliable source. Assuming that induction is conditionally reliable – that it gives true conclusions from true inputs – we thereby acquire animal knowledge that K is a reliable source. Once we have that knowledge, every piece of animal knowledge through K now converts into reflective knowledge.

The picture is actually somewhat more complicated than this for Sosa, since unlike Descartes, he admits several sources of knowledge, including perception, memory, introspection and induction (all of which figured in the track-record argument above). Reflective knowledge of the deliverances of perception would depend not just on animal knowledge that perception is reliable, but also on animal knowledge that the other sources are reliable. But the idea would remain that reflective knowledge is compounded out of animal knowledge in way that lets us break the skeptical impasse set up by (1) and (2).

Appraisal of the Coherentist and Two-levels Alternatives

We began with the concern that a certain kind of externalism makes it 'too easy' to know that our sources of knowledge are reliable. I raised the worry that the denial of this externalism, KR , would make knowledge impossible. We are now considering coherentism and the two-levels view as attempts to make knowledge possible, but not too easy. Have we succeeded?

Here is one doubt I have on this score. What is actually more demanding about the routes to knowledge recommended by the coherentist and the two-levels theorist? Let us consider a knower $S1$ as described by an externalist and a knower $S2$ as described by a coherentist. The first knower arrives at knowledge in the 'easy' way – he comes to know through introspection that perception has issued in his beliefs in p , q and r , through perception that these things are all true, and through inductive inference from this data that perception is a reliable source. 'Too easy', complains the coherentist, and suggests an alternative picture in which the perceptual beliefs are not really knowledge until the reliability of perception is known, and the reliability of perception not known until the truth of enough of its deliverances is known. We have knowledge of the particulars and the generals alike, but we do not have either until the moment at which we have the other. This is the more arduous path; now there is no easy knowledge.

But where is the extra toil? Might the knower as envisioned by the externalist and the knower as envisioned by the coherentist go through exactly the same steps, coming to believe all the same things at all the same times and concluding in the end that their belief-forming processes are reliable? If so, presumably either both of them know that their faculties are reliable or neither does. If $S2$ knows his faculties are reliable, so does $S1$. It cannot be said that $S1$ has followed too easy a route to arrive at knowledge, for he has followed the *same* route as the other knower in all non-epistemic respects.

To make my point another way, it looks as though the coherentist and the externalist-foundationalist agree in their answers to the first half and differ only in their answers to the second half of Senator Baker's question: what did he know and when did he know it? Items that are known sequentially for the one are known only at the end for the other. But it is hard to see how these differences in the times at which the title of 'knowledge' is bestowed can make *S2*'s knowledge either more genuinely knowledge or knowledge more laboriously obtained. If the coherentist's *S2* knows, so does the externalist's *S1*, even though the coherentist may have a more rigorous system of marking. This is a simple consequence of the principle of epistemic supervenience – that knowers alike in all non-epistemic respects must also be alike in what they know. (Similar points can be made about the knower as envisioned by the two-levels theorist with his expanded set of marks.)

But perhaps I am wrong in suggesting that the two knowers are in all non-epistemic respects the same. The coherentist might insist that the genuine knower must have a cognitive life that to some extent mirrors the story the coherentist tells. Perhaps the subject's degrees of belief must be apportioned to the accumulating evidence, the subject at first reposing only partial confidence in the various parts of his belief system until he sees how they all fit together. Perhaps the causal dependencies among the subject's beliefs must mirror the justificational dependencies charted out by the coherentist, the subject believing the particulars as strongly as he does partly because he believes the generals and vice versa. Perhaps the subject must be disposed to answer epistemological queries in coherentist fashion, not simply citing particular deliverances in support of general principles about reliability, but conversely as well, noting that his belief in the reliability of his faculties is part of what justifies him in accepting their deliverances. If such requirements are in force, then the road that must be followed by the coherentist knower is more difficult after all.

But perhaps it is *too* difficult. We are now requiring, in effect, that ordinary knowers are not really knowers unless they have sophistication matching that of the coherentist epistemologist. This leads Cohen to worry that his solution to the problem of easy knowledge may be psychologically unrealistic; it may also be part of what motivates Sosa's distinction between animal knowledge and reflective knowledge.

I am going to set this worry aside to consider a more fundamental misgiving. If externalism is problematic for reasons brought out by Vogel's gas-gauge example, so too, I believe, are coherentism and the two-levels view. So we should not abandon externalism in the hope that there are non-skeptical alternatives that avoid its problems.

The problem of too-easy knowledge is epitomized by Vogel's tale of the gas gauge. It seemed wrong for Roxanne to acquire knowledge of the reliability of her gauge and of her belief policies based on the gauge in the manner described. If it is wrong, it seems to me a wrong not rectified by imposing coherentist constraints on knowing. Let us suppose that initially upon looking at the gauge, Roxanne forms only a partial or cautious belief in 'The gauge reads *F*' and *F*'. Suppose that when she has accumulated many beliefs of that form, she tentatively draws the conclusion that her gauge is reliable. Suppose, however, that when she notices how the various

beliefs in her developing system reinforce each other, she believes all of them more firmly, ready to cite relevant others in support of any one. Suppose, in short, she behaves exactly as she should according to the coherentist's model of knowledge. Does she have knowledge of the reliability of her gauge, and of her own belief policies based on the gauge, that is not to be accorded to the original Roxanne? I think not. Anything that was troubling about the original example should trouble us still.

The same goes, I believe, for a two-levels Roxanne. We could tell the story again with Roxanne acquiring in the initial stages only animal knowledge concerning her gauge and its reliability and her own reliability in following it – animal knowledge which converts to reflective knowledge when there is enough of it in place. Would two-levels Roxanne gain knowledge that was denied to our two earlier Roxannes? I doubt it.

What is wrong with Roxanne's procedure? One cannot help but think it relevant that in the case of the gauge, there is a possibility of making independent checks (for example by inserting a stick in the tank or trying to add more gas when the needle says 'full') that Roxanne forgoes. In the case of an ultimate source of knowledge such as perception or memory, however, there is no such possibility of an independent check (and no hope for support except self-support). Upon noticing this, philosophers sometimes declare a crisis for the intellect. But perhaps another reaction is in order, as suggested by Nozick:

It seems plausible that philosophy should seek to uncover the deepest truths, to find explanatory or (if that is its aim) justificatory principles so deep that nothing else yields them, yet deep enough to subsume themselves. Reaching these should be a goal of philosophy, so when that situation occurs with some topic or area, instead of a crisis we should announce a triumph.

(1981: 137–8)

But this approach may also strike us as too easy, as brought out by Nozick's comparison of it with one senator's suggestion during the war in Vietnam that the USA should declare victory and then leave.

Conclusion

The problem of easy knowledge, as epitomized by the gas-gauge example, is not easy to solve. It remains with us if we abandon pure externalism for either a coherence or a two-levels view. I believe it is an outstanding challenge for epistemology to show how knowledge can be possible at all without being easy.

Notes

- 1 Among those who have reacted in this way are Fumerton (2000), Vogel (2000) and Cohen (2002).
- 2 I have been guided in what follows by Alston (1986a) and Bonjour (1992).

- 3 A third possibility as to the object of higher-order knowledge is precisely that one has first-order knowledge; the resulting internalism would be some version of the *KK* thesis.
- 4 The foregoing are all varieties of what has been called 'access internalism'. Another variety of internalism often encountered in recent literature is 'mental-state internalism', according to which justificatory factors must be mental, so that subjects alike in all mental ways must be alike in regard to what they are justified in believing. Access internalism of the 1bi and 2bi varieties entails mental-state internalism under certain plausible assumptions: if the only things knowable by reflection are mental states and necessary truths, then justificatory factors must be mental states, and they must confer justification in accordance with necessarily true principles. For more on these matters, see Ernest Sosa, 'Skepticism and the Internal/External Divide', in Greco and Sosa (1999: 145–57).
- 5 Compare p. 76: 'The internalist assumes that, merely by reflecting upon his own conscious state, he can ... find out, with respect to any possible belief he has, whether he is *justified* in having that belief.'
- 6 We shall consider a view to the contrary below, however, in connection with Thomas Reid.
- 7 We find confirmation in Chisholm (1977: 76): 'According to the traditional conception of "internal" epistemic justification, there is no *logical* connection between epistemic justification and truth.'
- 8 Insofar as the reliability need only be of the 'for the most part' variety, it is allowable that some of the premises tell of occasions on which I formed perceptual beliefs that I subsequently learned to be false.
- 9 Or of the argument's validity or inductive strength – a consideration that arises in inductive arguments for induction, but not in the track-record argument currently under consideration.
- 10 Or as Cohen says, it threatens to land us in 'the problem of the criterion'.
- 11 As Reid (1969) notes in Essay 6, Chapter 5, p. 631, 'Every kind of reasoning for the veracity of our faculties, amounts to no more than taking their testimony for their own veracity; and this we must do implicitly, until God give us new faculties to sit in judgment upon the old.'
- 12 I use the word 'first' to signal a relation of priority that is asymmetrical in the following sense: it is impossible that (*A* is prior to *B*, *B* is prior to *A*, and *A* and *B* both obtain). I do not necessarily mean by it to connote temporal priority, though that is perhaps the clearest case.
- 13 Although these passages make it highly natural to interpret Reid as affirming that the reliability of our faculties is epistemically basic, in 'Reid on the First Principles of Contingent Truths' I have made the case for an alternative interpretation in which Reid is instead an externalist who could countenance track-record arguments for reliability.
- 14 Indeed, in so far as 'my belief in *p* has a reliable source' factors into 'my belief in *p* has *K* as its source and *K* is reliable', *KR* seems to imply that any bit of knowledge depends on *two* further bits.
- 15 With one further proviso: we must not understand *KR* so that it would have as an instance 'S knows (*K* is reliable) through *K* only if S first knows *K* is reliable'. In other words, we must not allow *KR* to be rewritten as (1).
- 16 For another theory in which items acquire the status of knowledge by belonging to a system of mutually supporting elements, see Lehrer (1997).
- 17 For this reason, Cohen demurs at saying what is wrong with bootstrapping is that it involves self-support.
- 18 Reflective knowledge is governed by the following principles among others: epistemic ascent (if you know *p* and understand what knowledge is, then you know that you know

- p*), epistemic closure (if you know *p* and you know that *p* entails *q*, then you know *q*), and the principle of the criterion (if you know *p* and you understand what knowledge is, then you know that the source of your belief in *p* is reliable). See Sosa (1997b). The principles I have just mentioned are somewhat stronger than those Sosa formulates, but I believe he would accept them for reasons I explain in Van Cleve (2002).
- 19 Of course, if it were held that we cannot have reflective knowledge through *K* unless we first have *reflective* knowledge that *K* is reliable, skeptical consequences would again threaten. To avoid them, we would have to take a leaf from one of the two approaches canvassed above. Taking a leaf from Reid, we could say that there is *basic* reflective knowledge of the reliability of a source; or taking a leaf from the coherentist, we could say that reflective knowledge of reliability is not *prior* to reflective knowledge of the deliverances of a source, but coeval with it.
- 20 I am not sure that Sosa would endorse this biconditional, however, since it is not clear to me that all of his principles governing reflective knowledge are derivable if we adopt it.
- 21 We would presumably also need, for each of the original *cognitions*, a *cognitio* that its source was clear and distinct perception.