

ARTICLE

# Easy Knowledge, Circularity, and the Puzzle of Reliability Knowledge

Matthias Steup\*

University of Colorado, Boulder, USA

\*Email: [Matthias.steup@colorado.edu](mailto:Matthias.steup@colorado.edu)

## Abstract

According to externalist reliabilism and dogmatic foundationalism, it's possible to gain knowledge through a perceptual experience without being in a position to know that the experience is reliable. As a result, both of these views face the problem of making knowledge of perceptual reliability too easy, for they permit deducing perceptual reliability from particular perceptual experience without already knowing that these experiences are trustworthy. Ernest Sosa advocates a two-stage solution to the problem. At the first stage, a rich body of perceptual animal knowledge is acquired. At the second stage, perceptual knowledge becomes reflective after deducing perceptual reliability from the initial body of perceptual animal knowledge. I defend the alternative approach of rejecting both externalist reliabilism and dogmatic foundationalism. According to the alternative view, perceptual knowledge and knowledge of perceptual reliability require each other. Such a cognitive structure seems viciously circular. I propose that the appearance of vicious circularity dissipates when the relationship in question is viewed, not as one of temporal priority, but instead as synchronic mutual dependence. At a given time, one cannot have perceptual knowledge without knowledge of perceptual reliability, and vice versa. Such mutual dependence, I argue, is benign.

## 1. Introduction

The problem of easy knowledge raises, I believe, issues of crucial importance in epistemology. In my contribution to this volume in honor of Ernest Sosa, I will therefore discuss Chapter 10, “Easy Knowledge and the Criterion,” of his book *Reflective Knowledge*.

## 2. Fumerton on easy knowledge and externalism

Richard Fumerton brought the problem of easy knowledge into clear focus in his book *Metaepistemology and Skepticism* (Fumerton 1995). There, Fumerton objects to reliabilism as a paradigmatically externalist theory as follows:

If reliabilism is true, then introspection can give us justified beliefs that we are perceiving, and perception can give us justified beliefs that physical objects are present. The two reliable processes together can furnish a premise that, when

combined with others generated in a similar fashion, gives us inductive justification for believing that perception is reliable.<sup>1</sup>

Fumerton's point is that reliabilism is problematic because it makes it too easy for us to come to know that perception is reliable. Why too easy? Fumerton says: because the reasoning employed is circular. Perception is used to confirm the reliability of perception. Here is a choice quote:

If as a philosopher I start wondering whether perceptual beliefs are accurate reflections of the way the world really is, I would not dream of using perception to resolve my doubt.<sup>2</sup>

Fumerton, then, rejects the easy way in which reliabilism allows for knowledge of perceptual reliability because it is inflicted with an unacceptable kind of source circularity. I believe that Fumerton's objection to reliabilism is deep and important, but I don't agree with his diagnosis of why reliability knowledge cannot be acquired in the easy way reliabilism permits.

### 3. Cohen on easy knowledge and the KR principle

In "Basic Knowledge and the Problem of Easy Knowledge," Stewart Cohen adds two important insights to our understanding of the issue.<sup>3</sup> First, the problem of easy knowledge afflicts not only externalist reliabilism, but also internalist theories such as James Pryor's dogmatism and Michael Huemer's phenomenal conservatism.<sup>4</sup> Second, the bootstrapping-type of reasoning Fumerton rejects as defective is just one type of easy knowledge. Another type is unacceptably easy knowledge that one's current perceptual experiences are not deceptive.

What licenses both bootstrapped reliability knowledge and easy knowledge of non-deception is the rejection of what Cohen calls the "KR principle":

KR A potential knowledge source, S, can yield K for you → you know that S is reliable.<sup>5</sup>

It is important to distinguish between the key idea underlying KR and specific versions of the principle. It seems to me that making the key idea explicit does not require articulating the principle in terms of knowledge. Those who endorse KR typically do so from an internalist point of view, according to which a belief source is a knowledge source by virtue of being a justification source. Thus we can move from KR to:

KR<sub>a</sub> A belief source, S, is a justification source for you → you know that S is reliable.

Furthermore, we might worry that KR is too demanding if the kind of knowledge that figures in KR's consequent is fully explicit, thus entailing occurrent belief. In ordinary

<sup>1</sup>Fumerton (1995: 176). See also Vogel (2000).

<sup>2</sup>Fumerton (1995: 177).

<sup>3</sup>Cohen (2002).

<sup>4</sup>See Pryor (2000) and Huemer (2001). Cohen refers to such views as "evidentialist foundationalism." It would be a mistake, however, to think the problem of easy knowledge arises for internalist foundationalism as such. It arises only for versions of internalist foundationalism that deny what Cohen calls the "KR principle." I will return to this point further below.

<sup>5</sup>Cohen (2002: 309).

life, perceptual experiences result in spontaneously formed perceptual beliefs, without being accompanied by occurrent beliefs about the reliability of these experiences. Therefore,  $KR_a$  implies that, when perceptual beliefs are spontaneously formed, perceptual experiences are not a justification source. The following modification of the consequent avoids this problem:

$KR_b$  A belief source,  $S$ , is a justification source for you  $\rightarrow$  you are in a position to know that  $S$  is reliable.

But what is it to be in a position to know? Different theories will offer opposing answers. From the internalist approach that motivates  $KR$ , what puts one into a position to have knowledge is the possession of evidence. Thus we arrive at:

$KR_c$  A belief source,  $S$ , is a justification source for you  $\rightarrow$  you have evidence for taking  $S$  to be reliable.

A further question of interpretation arises when we ask exactly what the intended relation is between the antecedent and the consequent. On one approach, it is an entailment relation: necessarily, if the antecedent is true, then the consequent is true. I believe that interpretation fails to capture the key thought that motivates  $KR$ . When epistemologists advocate  $KR$ , the question they are concerned with is: When a belief source is a justification source for you, by virtue of what does it enjoy this status? What *makes* it a source of justification for you? Epistemologists who endorse the key thought motivating  $KR$ -type principles would answer this question thus: whenever a belief source,  $S$ , is a justification source for you, it has that status because of your reasons or evidence in support of the proposition that  $S$  is reliable. This evidence is what makes  $S$  a justification source for you.<sup>6</sup> Accordingly, I will understand  $KR_c$  as follows:

$KR_d$  A belief source,  $S$ , is a justification source for you  $\rightarrow$  (i) you have evidence  $E_R$  for taking  $S$  to be reliable and (ii)  $E_R$  is what makes  $S$  a justification source for you.

In what follows, I will understand the  $KR$  principle in just this way.<sup>7</sup> Obviously, there are additional ways of interpreting  $KR$  resulting from alternative construals of the epistemic statuses that feature in the principle's antecedent and consequent. Exploring this territory more deeply would be a topic for a different paper.

Is  $KR_d$  sufficiently plausible to spell trouble for theories that reject it? I believe it is. What motivates the principle is the basic thought that it's irrational to rely on a source without having reasons to trust it.<sup>8</sup> Both reliabilism and dogmatism reject that thought.

<sup>6</sup>When I talk of having evidence that source  $S$  is reliable, what I mean is having reasons, or having justification, for taking  $S$  to be reliable. The concept of justification I'm working with is internalist: what makes your beliefs justified is what you have to go on, that is, the way things seem to you from your own point of view.

<sup>7</sup>It might be argued that, when choosing a suitable version of  $KR$ , pressing the easy knowledge objection raises a dilemma: a particular version of  $KR$  is either so demanding that it isn't plausible, or it is plausible but then too weak for generating the problem of easy knowledge. My response is that  $KR_d$  is weak enough to be plausible: it does not burden subjects with excessive demands for epistemic reflection. At the same time, since reliabilists and dogmatists would deny  $KR_d$ , the principle is strong enough for generating the easy knowledge problem.

<sup>8</sup>When we apply  $KR$ -type principles to testimony, we are looking at the debate between reductionists and anti-reductionists. The former endorse, the latter reject, testimonial versions of  $KR$ . Since I believe that  $KR$ ,

According to these views, for a belief source to give you knowledge, you needn't know, nor be in a position to know, nor have any kind of justification at all for believing, that the source is reliable. That's why they face the problem of easy knowledge.

#### 4. Bad deductions: easy knowledge of non-deception

Fumerton points out that reliabilism allows for knowledge of reliability that is suspiciously easy. Cohen expands Fumerton's point by showing that any theory that denies KR suffers from this problem, and from the additional problem of permitting surprisingly easy knowledge of non-deception. Let's see how, according to Cohen, denying KR allows for the latter kind of easy knowledge. The example Cohen describes features his son who wants to know whether the red-looking table he and his dad might purchase is really red. Reliabilism and dogmatism allow resolving this doubt in the following way. Suppose, first of all, the situation is entirely normal. The table is illuminated by daylight and looks red to Cohen and his son because their color vision works reliably and the table is in fact red. Now suppose reliabilism is true. Then Cohen and his son know the table is red because their color vision is reliable. Suppose dogmatism is true. Then Cohen and his son know the table is red because it looks red to them. In each case, they come to know that the table is red without needing any justification for believing that their visual experiences are reliable.<sup>9</sup> Equipped with such knowledge, they can deploy the following deduction:

##### A Bad Deduction

- (1) The table is red.
- (2) (The table is red)  $\rightarrow$   $\sim$ (The table is white and illuminated by red lights)
- (3)  $\sim$ (The table is white and illuminated by red lights)

What are the grounds for calling this deduction 'bad'? After all, logically there is nothing wrong with it. And, *ex hypothesi*, the premises are true. The deduction, then, is both valid and sound. Therefore, the reason for rejecting it as bad is not that the argument is logically flawed or has a false premise. Rather, the point is that the deduction fails to do what arguments are supposed to do, namely either track pre-existing or generate new justification for accepting the conclusion. When an argument is presented to resolve a doubt or justify a claim, its job is to answer a certain question, namely: How do we know the claim in question? For a proffered argument to answer this question, its premises must track the evidence in support of the claim that has been questioned. The argument from (1) and (2) to (3), then, is supposed to identify the evidence on the basis of which Cohen and his son know that the red-looking table isn't white with red lights shining at it. If Cohen were to offer this argument as a response to his son's worry, he would be telling his son that knowledge of (3) comes about by deducing it from (1) and (2). The point of Cohen's example is that, when his son doesn't buy this argument, he recognizes that (3) is *not* known, and cannot come to be known, by deducing it from (1) and (2). Therefore, the argument fails to generate justification for its conclusion. That is the sense in which it is a 'bad' deduction.

---

understood as KR<sub>d</sub>, applies to all justification sources, I side with reductionism. For relevant literature, see, for example, Fricker (1994), Lackey (2008), and the essays in Lackey and Sosa (2006).

<sup>9</sup>Of course, further conditions are at play as well: Cohen and his son know that the table is red because it's true that that the table is red, and because they don't have any defeaters, and because they are not in a Gettier-type situation.

Let's focus on easy knowledge reasoning in the hands of dogmatists. According to dogmatism, a perceptual experience as of *p* can *by itself* justify the belief that *p*. Its justificational force does not derive from justification for taking it to be reliable, or from justification for rejecting deceptive scenarios. So, if dogmatism is true, when the table looks red to Cohen and Cohen has no defeaters, Cohen thereby comes to know, solely on the basis of the table's *looking* red to him, that the table *is* red. Consequently, the Bad Deduction reduces in essence to the following bit of reasoning:

*No Red Light:* I know the table isn't white with red lights shining at it because it looks red to me.

In the same vein, dogmatism licenses arguments like the following:

*No Painted Mules:* I know the animals in the pen are not cleverly painted mules because they look like zebras.

*No Hallucination:* I know I'm not hallucinating an oasis because I'm having a visual experience of palm trees and a pond.

*Non-Envatment:* I know I'm not a BIV because I have hand experiences.

It's difficult to not have the feeling that these are epistemically defective arguments, not because of any logical defects but rather because they fail to track the grounds on which the propositions in questions are actually known. That is what makes them bad deductions. Dogmatism and reliabilism license bad deductions. Therefore, they face the problem of easy knowledge.<sup>10</sup>

## 5. Why bad deductions are bad: transmission failure

Martin Davies and Crispin Wright have argued that bad deductions suffer from transmission failure: they fail to transmit justification from the premises to the conclusion.<sup>11</sup> *No Red Light* fails in this regard because perceptual knowledge is not as easily acquired as dogmatism would have it. According to dogmatism, the table's looking red can by itself give you knowledge that the table is red. The competing, more demanding view says that the table's looking red can give you knowledge that the table is red only if you have, *to begin with*, justification for rejecting the white table and red lights alternative. Knowing the argument's first premise (that the table is red) requires prior justification for believing the argument's conclusion. *No Red Light*, then, fails to transmit justification, and thus fails to generate knowledge, because it reverses the direction of justification.

Consider *Non-Envatment*. If hand experiences are a source of hand knowledge for me, I must already have justification for rejecting any skeptical scenario in which I have hand experiences without having hands. Therefore, justification for the premise

<sup>10</sup>It's important to note the following: the criticism is not that reliabilism and dogmatism don't allow for sound or non-easy knowledge of reliability and non-deception. According to reliabilism, knowledge of reliability is obtained whenever a belief such as "x is reliable" is true and produced by an undefeated reliable process. According to dogmatism, reliability knowledge is acquired if one has an undefeated, non-misleading seeming of reliability. Analogous points apply to knowledge of non-deception. Rather, the criticism is that reliabilism and dogmatism allow, *in addition*, for a clearly defective way of obtaining knowledge of reliability and non-deception.

<sup>11</sup>See Davies (2004) and Wright (2004). See also Brown (2005).

“I have hands” does not transmit to the conclusion “I’m not a BIV.” Instead, I know that premise only if I have, to begin with, justification for rejecting the BIV scenario. Analogous points apply to *No Painted Mules* and *No Hallucination*.

Regarding *No Red Light*, Sosa puts the point very nicely in terms of what one must presuppose if one’s color experiences are to give one knowledge. He points out that, even if KR is false,

it may still be that in order for you to know that the [table] is red it is not enough to base your belief that it is red simply on its looking red. Perhaps your belief must be based on that visual appearance along with a presupposition that your situation is relevantly unproblematic.<sup>12</sup>

Next, Sosa emphasizes that this presupposition must not be arbitrary. It must be epistemically appropriate or justified. That’s the rub, for then it follows that

you cannot then posteriorly acquire this appropriateness (amounting perhaps to knowledge), based on the knowledge that presupposes it, namely the knowledge that the [table] is then red.<sup>13</sup>

This seems essentially right to me, although I myself would prefer to state the relevant point in terms of evidence: for visual experience to have justificational force for you, you must have evidence that the present conditions don’t undermine the trustworthiness of your experience.

The key thought motivating the rejection of easy knowledge is that of *epistemic priority*. Perceptual knowledge has a directional structure: having justification for rejecting reliability-undermining scenarios is epistemically prior to, or a presupposition of, the acquisition of perceptual knowledge. Therefore, we cannot first acquire perceptual knowledge and then deduce that reliability-undermining deceptive scenarios don’t obtain.

## 6. Attempts to defuse the easy knowledge challenge

Since dogmatists deny KR-type principles, they face the problem of easy knowledge. To defuse it, they have two options.<sup>14</sup> First, they could claim that bad deductions aren’t really bad. Second, they could claim that, initial impressions notwithstanding, dogmatism does not actually license bad deductions. Since the bad deductions listed above are paradigmatic examples of begging the question, the first option doesn’t seem very attractive. James Pryor and Peter Markie have suggested a hybrid strategy. In a context of doubt, the bad deductions are indeed bad because they beg the question. But that’s merely a dialectical defect, a defect that is limited to the sphere of rational persuasion. Dogmatists, the response goes, can consistently concede this defect and nevertheless

<sup>12</sup>Sosa (2009: 219).

<sup>13</sup>Sosa (2009: 220).

<sup>14</sup>How might reliabilists respond to the easy knowledge argument? An initial move might be to say that easy knowledge reasoning is unreliable. Therefore, it is false that reliabilism licenses such reasoning as knowledge generating. The problem with this move is that, while easy knowledge reasoning is unreliable when employed by subjects who are actually deceived, it’s unclear why it should be unreliable in the hands of subjects who are not deceived. For example, when Cohen responds to his son’s doubts with *No Red Lights*, he reasons from true premises to a true conclusion. It’s not easy to see why such reasoning should be unreliable.

insist that the allegedly defective arguments can generate knowledge of their conclusions.<sup>15</sup>

This response assumes that an argument's begging the question is one thing and its capacity to produce knowledge another. I don't buy this assumption. It seems to me that, when a proffered argument is rejected as question begging, the point is that the rejected argument fails to explain how the claim in question could be known. For the response to go through, the defect in question would have to be merely psychological: a failure to persuade someone who harbors doubt about the argument's conclusion. But when we consider the full range of bad deductions and find them wanting, the thought is not that what's wrong with these arguments is merely that they are impotent means of persuasion and thus ill suited to placate pesky doubters. Rather, the worry is that they all fail as accounts of how the conclusions in question are actually known.<sup>16</sup>

## 7. The puzzle of reliability knowledge

Views that reject KR-type principles are problematic because they allow for knowledge that's too easy. But views that accept such principles are also problematic. To see why, let us focus on perception. Applied to perception, KR principles tell us that perceptual knowledge requires prior knowledge (or at least evidence) of perceptual reliability. The problem is that there is a second principle, at least as plausible as KR if not more, namely RK: knowledge of perceptual reliability requires prior perceptual knowledge. The point is obvious enough: unless it's possible to deduce perceptual reliability a priori, you have to use your perceptual faculties to find out whether perception is reliable. Now, it would seem that, if both KR and RK are true, neither perceptual knowledge nor knowledge of perceptual reliability is possible. To have one, you must have the other one first: a catch 22 situation. The puzzle, then, arises because KR and RK are both individually plausible while their conjunction seems to form a vicious circle. Spelled out, the puzzle takes the following form:

KR Your perceptual experiences are a knowledge source for you → you have justification for taking them to be reliable.<sup>17</sup>

RK Having justification for taking your perceptual experiences to be reliable → your perceptual experiences are a knowledge source for you.

NP If KR and RK are both true, then neither perceptual knowledge nor knowledge of perceptual reliability is possible.<sup>18</sup>

Therefore:

BN Neither perceptual knowledge nor knowledge of perceptual reliability is possible.<sup>19</sup>

<sup>15</sup>See Pryor (2004) and Markie (2005). White (2006) rejects dogmatism but on p. 29 expresses some degree of sympathy for the line of argument Pryor and Markie deploy.

<sup>16</sup>For further discussion, see Cohen's rather effective response to Markie in Cohen (2005).

<sup>17</sup>This version of KR differs from the version on which I settled earlier: KR<sub>d</sub>. For the purpose of easy exposition I'm keeping the principle simple here, but I do understand it as KR<sub>d</sub> applied to perception.

<sup>18</sup>"BN" stands for "bad news" and "NP" for "not possible."

<sup>19</sup>In his 2000, Pryor argues that the standard deception-type of skeptical argument – appealing, for example, to the possibility of being a BIV – is not the strongest version of skepticism. It seems to me that the alternative, more serious argument he envisions is best construed as the puzzle I'm discussing here. And while I would agree that the puzzle poses a formidable challenge, I would, however, reject Pryor's suggestion that standard deception arguments (of the evil demon or BIV type) are less of a challenge.

Here are the options for escaping the puzzle:

1. Reject KR. That's what externalists and dogmatists recommend. To make this response palatable, they must find a way to defuse the charge of easy knowledge.
2. Reject RK. For this move to be workable, there must be an *a priori* argument for the reliability of our perceptual faculties.<sup>20</sup> It's not easy to see what that argument might be.
3. Reject NP. Advocates of this option face the challenge of explaining why the conjunction of KR and RK does not constitute a vicious circle.

I believe that the third of these options is the most promising, and I'll provide some support for this view at the later stages of this paper.

### 8. Dogmatism, conservatism, and credentialism

The puzzle of reliability knowledge is closely related to three main options to consider when we ask what it takes for a perceptual experience as of *p* to be a reason – a source of justification – for believing *p* and thereby to be a source of knowledge of *p*.<sup>21</sup> We have already come across the first option:

**Dogmatism:** A perceptual experience as of *p*, *Ep*, is by itself a reason for believing *p* *no matter what*. The mere having of the experience is sufficient. For it to be a knowledge source, it must not meet any qualifying condition. Consequently, even if you have no evidence at all that *Ep* is reliable, *Ep* is nevertheless a reason for you to believe *p*. In fact, it is such a reason even if you have evidence that *Ep* is unreliable.<sup>22</sup>

Conservatives don't agree that perceptual experiences are a knowledge source no matter what. They say the following:

**Conservatism:** A perceptual experience as of *p* is a reason for you to believe *p* provided you have no reason to believe that the experience is unreliable.<sup>23</sup>

<sup>20</sup>Cohen considers this avenue in his 2002 and defends it in his 2010 (Cohen 2002, 2010).

<sup>21</sup>There are of course further, externalist options. On a simple reliabilist view without bells and whistles, perceptual experience is a source of justification if and only if it is reliable. According to approaches that combine internalist and externalist elements, *de facto* reliability is merely a necessary condition: perceptual experience is a source of justification only if it is in fact reliable. For such a view, see, for example, Brogaard (2013).

<sup>22</sup>This is not to say that believing *p* on the basis of *Ep* is justified even in the presence of evidence that *Ep* is unreliable. Perceptual justification is defeasible, and evidence of unreliability is a defeater. Thus, the claim in question is not that *Ep* is sufficient for *being justified in believing p*. Rather, the claim is that *Ep* is sufficient for *having justification for believing p*. Dogmatism, then, is a thesis about propositional, not doxastic, justification. Here are two choice quotes expressing the dogmatic view. Pryor: "My view is that whenever you have an experience as of *p*, you thereby have immediate *prima facie* justification for believing *p*" (Pryor 2000: 536. Cf. p. 532). Huemer: "If it seems to *S* as if *P*, then *S* thereby has at least *prima facie* justification for believing that *P*" (Huemer 2001: 99). For a defense of dogmatism, see also Tucker (2010).

<sup>23</sup>Conservatism in epistemology comes in at least three forms: (i) as a claim about what's sufficient for a belief to be justified, (ii) as an approach to offering epistemic guidance to belief formation, and (iii) as a claim about when a belief source is a source of justification. The following passage in Chisholm (1980: 551) is an example of (i): "Anything we find ourselves believing may be said to have some presumption in its favor." Another example of (i) is Harman (1986). However, he also advocates (ii). For another

Conservatives agree with dogmatists that, for Ep to be a reason for you believing p, you need not have evidence that Ep is reliable. But they reject the following implication of dogmatism: Ep is a reason for believing p even if you have evidence that Ep is unreliable. Advocates of KR go even further. They reject the conservative claim that Ep is a reason for believing p even if you don't have evidence that Ep is reliable. Instead, they hold that, for Ep to be a knowledge source for you, you need to have positive evidence that Ep is reliable. Thus we arrive at the third view:

Credentialism: A perceptual experience as of p is a reason for you to believe p only if you have evidence, ideally strong enough for knowledge, that the experience is reliable.<sup>24</sup>

I label this view 'credentialism' because it tells us that a perceptual experience is a source of justification only if it is properly credentialed by evidence of its reliability.

Here is how these views relate to the puzzle of reliability knowledge. Dogmatists reject KR. They must somehow find a way to defuse the charge of easy knowledge. Credentialists accept KR. If they accept RK as well, they must explain why the conjunction of KR and RK is not an epistemic debacle. Vis-à-vis these two positions, conservatives enjoy a couple of advantages. Unlike dogmatists, they do not claim, implausibly, that perceptual experience is a justification source no matter what. Rather, they hold that, when a perceptual experience is undermined by evidence of reliability, it ceases to be a justification source. Unlike credentialists, they do not go as far as saying that a perceptual experience is a source of justification for you only if you have evidence that it is reliable. Conservatives, then, do not accept KR principles, and therefore avoid the circularity problem that results from conjoining KR and RK.<sup>25</sup> However, these advantages notwithstanding, since conservatives reject KR, they face the problem of easy knowledge no less than dogmatists do.<sup>26</sup>

## 9. Two kinds of basicness

Cohen calls 'basic knowledge' the kind of knowledge that is possible if KR is false: knowledge based on a body of evidence that doesn't include any evidence that its source

---

example of (ii), see Lycan (1998, 2013). (iii), which is clearly a theoretical option, has not, as far as I can tell, been explicitly defended. Huemer advocates what he calls "the principle of phenomenal conservatism," but the principle actually articulates dogmatism, not conservatism as I understand it here. Whether his theory is conservative in some other sense is an interesting question, but one that cannot be addressed in this paper. In addition to (i)–(iii), there is now in the literature a new and different use of the term 'conservatism,' according to which conservatism is to be identified with the rejection of dogmatism. According to this terminology, dogmatism is a form of liberalism, which is opposed by more demanding view that qualify as conservatism. This alternative (and in my view not entirely fortuitous) use of the term 'conservative' originates with Pryor (2004). See also Silins (2007) and Neta (2010).

<sup>24</sup>In his 2002, Cohen refers to this view as 'holism,' a label associated with Pierre Duhem and W.V.O. Quine. In his 1951, Quine describes the view thus: "Our statements about the external world face the tribunal of sense experience not individually but only as a corporate body" (Quine 1951: 41). While Quine was concerned primarily with meaning, the holistic approach can be extended to justification as well, resulting in the view that a perceptual belief's justification is conferred on it not by a single sense experience but instead by the subject's entire belief system. A more recent example is BonJour (1985). The holistic view I have defended in Steup (2004, 2018) differs crucially from BonJour's in the following respect: Whereas according to BonJour (1985) empirical justification derives from a web of beliefs, on my view, it derives from a web of seemings.

<sup>25</sup>The circularity problem arises when the presence of reliability evidence is required. When merely the absence of evidence of unreliability is required, the circularity problem is avoided.

<sup>26</sup>For further discussion of conservatism, see Christensen (1994), Vahid (2004), and Steup (2018).

is reliable. For example, when a table's looking red to me *by itself* – unaided by any evidence of reliability – gives me knowledge that the table is red, then my knowledge that the table is red is basic in Cohen's sense. Views that deny KR are, in Cohen's terminology, 'BKS theories': theories with a basic knowledge structure. This terminology suggests that traditional foundationalism, since it postulates the existence of basic knowledge, is a BKS theory. But it would be a mistake to think that basic knowledge as Cohen defines it is the same as the kind of basic knowledge that foundationalists say must exist if there is any knowledge at all. Rather, we must, as Sosa suggests, distinguish between the following two kinds of basicness:

- (i)  $K_p$  is *inferentially basic* iff  $K_p$  is not based on inference, so its sources are not inferential.
- (ii)  $K_p$  is *hierarchically basic* iff  $K_p$  is unaccompanied by knowledge that its basis (or, more generally, its source) is reliable.

Sosa adds: "Thus, a bit of knowledge could be inferentially basic without being hierarchically basic, and a bit of knowledge could be hierarchically basic without being inferentially basic."<sup>27</sup>

Sosa's distinction between inferential and hierarchical basicness is important. It will be worth our while to focus on it. In addition to the two possibilities Sosa mentioned, there are two further ones. Here is a list of all four:

1. I-basic & H-basic
2. I-basic &  $\sim$ (H-basic)
3.  $\sim$ (I-basic) & H-basic
4.  $\sim$ (I-basic &  $\sim$ (H-basic))

I will review all four cases. But before I do so, I wish to provide a bit more theoretical background about the two types of basicness at play.

## 10. Justification and meta-justification

The kind of basic justification dogmatists have in mind is possible only if the KR principle is false. Pryor calls such justification 'immediate.' I do not think using the label 'immediate' is entirely fortuitous because it blurs the distinction between the two types of basicness we must separate. If your belief that the table is before you is red is justified *solely* by its looking red to you, then, I agree, it is properly described as being immediately justified. But a belief is also properly described as immediately justified if its justification does not come from any additional belief and thus is non-inferentially justified. These two types of immediacy must be distinguished because, as Sosa points out, non-inferential justification need not coincide with the kind of immediacy dogmatists have in mind.

The label 'non-hierarchical' is not entirely free of ambiguity either. Inference creates a hierarchy of premises and conclusions. A belief that's justified via inference from premise-beliefs is hierarchically justified in a way that non-inferentially justified beliefs are not. Hence, it's necessary to spell out in exactly which sense the kind of basicness that's incompatible with the KR principles is non-hierarchical. Here is how I think this can be done. There is a kind of hierarchical structure – marking two different types of justification – that is the central concern in the dispute over the KR principle, namely 1st order vs. meta or 2nd order justification. For example, a table's looking red gives you

<sup>27</sup>Sosa (2009: 228).

1st order justification for believing it is red. Evidence indicating that your visual experience is reliable gives you 2nd order justification or meta-justification for believing that the table is red. In general terms, meta-justification is justification for believing that the things that give you 1st order justification are reliable or trustworthy. KR principles express the thought that there is no justification, and thus no knowledge, without meta-justification. Opponents of the KR principle believe that, if there is any knowledge at all, there must be justification without meta-justification.

Using the distinction between 1st order and 2nd order justification, we may now describe the two types of basicness as follows:

Inferential basicness: the kind of *non-doxastic justification* that traditional foundationalists say must exist if there is any knowledge at all.

Hierarchical basicness: the kind of *1st order justification without meta-justification* that dogmatists say must exist if there is any knowledge at all.

Here is what's important about this distinction: inferential basicness does not require hierarchical basicness. It doesn't because, as we will see momentarily, it's possible that I have non-doxastic perceptual justification for believing *p* and, at the same time, also have meta-justification for believing *p*.

### 11. I-basic and H-basic

Knowledge that is both inferentially and hierarchically basic has these two features: it does not result from inference from another belief, and its source does not include any kind of meta-justification: no evidence indicating that its source is reliable. Here are a couple of examples:

- (i) You know that the table before you is red solely because it looks red to you. At the same time, you are not in a position to know that, under the circumstances, looking-red is a reliable indicator of being-red because you have zero evidence that the table is not illuminated by red light.
- (ii) You know that the animals in the pen are zebras solely because they look like zebras to you. At the same time, you are not in a position to know that, under the circumstances, looking-like-a-zebra is a reliable indicator of being-a-zebra because you have no evidence allowing you to know that the animals in the pen are not cleverly disguised mules.

Dogmatism should not be understood as the view that these two cases describe our ordinary situation when a table looks red to us or animals in the zoo look like zebras. Normally, the presence or absence of red light is perfectly obvious. And normally we do have very good evidence for rejecting the cleverly-disguised-mule scenario. In ordinary life, therefore, it's not so easy to find counterexamples to KR. Dogmatism does, however, allow for the *possibility* of cases (i) and (ii) above. The view allows for this possibility precisely because it denies KR. If KR is false, if perceptual knowledge does not require evidence of perceptual reliability, then a perceptual experience that *p*, if in fact reliable and resulting in a true belief that *p*, is sufficient for coming to know that *p* is true. From knowledge of *p*, it can then be deduced that any deceptive  $\sim p$  situation does not obtain. According to the objection from easy knowledge, coming to know one isn't deceived isn't quite as easy as that.

## 12. I-basic and $\sim$ (H-basic)

I consider this category the main take-away of Sosa's distinction. Its significance is that foundationalism is not the same as dogmatism. Why? Consider the standard definition of the kind of belief foundationalists take to be basic:

B is a justified basic belief iff (i) B is justified, but (ii) B's justification does not come from any further belief, B\*.<sup>28</sup>

A belief that is non-inferentially justified in this way has a type of justification that is non-doxastic, one that does not come from further beliefs. So, the answer to the question,

Why is non-inferential justification not necessarily non-hierarchical justification?

is: because evidence of perceptual reliability need not come in the form of *beliefs* about perceptual reliability.

Imagine yourself in front of a red-looking table. You are going to have introspective and perceptual evidence regarding the present circumstances. This evidence, were you to consider it, would tell you the following: You have not taken any vision-distorting drugs, you don't suffer from a brain tumor causing perceptual malfunction, the table is illuminated by ordinary daylight coming through a large window. Your memory, if you were to consult it, would tell you that in the past, your color vision has, under such conditions, been nearly always accurate. This is a rich body of evidence for the belief that, under the present circumstances, the table's looking red to you is reliable. It is evidence that you can trust your visual experience. The crucial point is this: if we understand KR as requiring not knowledge but merely evidence of reliability, then KR does not require that, for the table's looking red to be a knowledge source for you, you must form explicit beliefs on the basis of the reliability evidence I just described. *Having* this body of evidence is sufficient for satisfying what KR demands.

The language of seemings is useful to press this point. Consider again the example of the red-looking table. You have introspective seemings about the content of your visual experience. You have perceptual seemings about the size of the table, about how far away it is from you, and about the lighting conditions. Finally, you have memorial seemings that, in the past, perceptual experiences under such conditions have been highly accurate. Suppose further that you don't form any beliefs about all of that. Thus your belief that the table is red is based just on seemings, not on any further beliefs. Your red table belief, therefore, is a non-inferential belief – the kind of belief foundationalists consider basic. Nevertheless, your epistemic situation satisfies a principle like KR<sub>D</sub>. Therefore, foundationalists are free to endorse KR<sub>D</sub>.

## 13. $\sim$ (I-basic) and H-basic

Sosa suggests that “a bit of knowledge could be hierarchically basic without being inferentially basic.” What would be an example of such knowledge? Let's consider a candidate for this status: an inferential belief that is hierarchically basic. Assume you infer and thereby come to know that your car's tank is half-full because the needle of the gas gauge points to ½. Can this bit of knowledge be hierarchically basic? If it is, then its evidential basis does not involve any meta-justification: it involves no reasons for

<sup>28</sup>See, for example, Feldman (2003: 50).

believing that the gas gauge is reliable or trustworthy. The question, then, comes to this: Can you come to know

H The tank is half-full

by inferring it from

G The gas gauge needle points to  $\frac{1}{2}$

without having any evidence for

R The gas gauge is reliable.

Gas-gauge dogmatism rejects the KR principle applied to gas gauges:

$KR_{\text{gas-gauge}}$  Your car's gas gauge is a knowledge source for you  $\rightarrow$  you have evidence that the gas gauge is reliable

By itself, gas-gauge dogmatism does not suffice to yield the result that your knowledge of the tank's being half-full is non-hierarchical. To secure the non-hierarchical nature of your knowledge, we also need a kind of inference dogmatism. This view rejects the KR principle applied to the inference in question:

$KR_{\text{inf}}$  Inference based on the gas gauge's needle-position is a knowledge source for you  $\rightarrow$  you have evidence that this type of inference is reliable.

For your knowledge of H to fit the category under examination, a radical form of dogmatism would have to be true: the kind of dogmatism that rejects both  $KR_{\text{gas-gauge}}$  and  $KR_{\text{inf}}$ . If one thinks that dogmatism of this kind is implausible, one might doubt that there really is inferential knowledge that is hierarchically basic.

#### 14. $\sim$ (I-basic) and $\sim$ (H-basic)

Suppose I come to know that the tank is half-full because I infer it from (i) the gas gauge's needle position and (ii) from my justified belief that my car's gas gauge is reliable. Given (i) and (ii), my knowledge is neither inferentially basic nor hierarchically basic. It seems unproblematic that knowledge of this type exists. What is controversial is its scope. In *The Structure of Empirical Knowledge*, Laurence Bonjour argued that all empirical knowledge is like that.<sup>29</sup> The view Bonjour defended in this book is extreme: justified perceptual beliefs are always both inferentially and hierarchically justified because they receive their justification from meta-beliefs about perceptual reliability.

Epistemologists who like the KR principle need not agree with Bonjour's doxasticism: his view that meta-justification must come from beliefs. If they endorse KR but reject Bonjour's doxastic coherentism, they thereby take a foundationalist position according to which perceptual knowledge does not require meta-justificatory beliefs about perceptual reliability. It merely requires evidence of perceptual reliability, giving us justification for believing that perception is reliable were we to reflect on the matter. On this view, all perceptual knowledge is hierarchical, but not all of it is inferential. Dogmatists advocate a

<sup>29</sup>Bonjour (1985).

type of foundationalism that is more radical than that: dogmatic foundationalism, according to which some perceptual knowledge is neither inferential nor hierarchical.

### 15. The problem of non-existing knowledge sources

Although Sosa is on balance friendly to KR, he raises two problems for it. First, some instances of knowledge might not come from any source. Second, in many cases, what source a bit of knowledge comes from might be entirely hidden. I consider the first challenge in this section and the second in the next.

As examples of knowledge without a source, Sosa mentions knowledge of the following two propositions: “I exist” and “something exists or not.”

Here, two questions must be asked:

1. Assuming there is source-less knowledge, how much of a problem would that be for advocates of the KR principle?
2. Is there source-less knowledge?

In response to the first question, the following two points seem relevant. First, instances of source-less knowledge do not seem to be counterexamples to the KR principles because they are instances of a type of knowledge that falls outside the scope of KR. If knowledge of *p* is source-less knowledge, it doesn't come from a source whose reliability, according to KR, needs to be known. Second, if there is source-less knowledge, then there is a type of knowledge that is unaffected by the twin problems of easy knowledge and vicious circularity. However, even if there was such a type of knowledge, a lot of our knowledge would still come from sources. If we wish to explain how such knowledge is possible, we need to find a solution to these problems.

In response to the second question, I believe skepticism is warranted. If we search for the sources of knowledge of our own existence, there are certainly good candidates that suggest themselves. One is perception. I can see and feel my body. Another one is introspection. I'm introspectively aware of my current thoughts. And, as Descartes pointed out, it's difficult to see how I could be having these thoughts if I didn't exist. It would seem, therefore, that knowledge of my own existence is a plausible instance of the KR principle: it has a source, and I can know that this source is reliable (and indeed infallible). Analogous reasoning applies to axiomatic a priori knowledge. I know “Something exists or not” through a priori intuition, and through a priori intuition I can know that this proposition cannot be false. Again, we seem to have a case not only of knowledge coming from a source, but knowledge that fits the KR principle. For, if intuition tells me that *p* cannot be false, then I have excellent evidence for thinking that my belief that *p* comes from a reliable source.

Perhaps, when Sosa considers the possibility that some knowledge is without a source, what he has in mind is knowledge without a *rational* source. The worry might go as follows. My belief in my own existence has its source way back in my distant past. On that approach, the source is identified with the causal process that originally produced my belief. It seems clear that, if we think of my belief's source in that way, I do not have cognitive access to its source. It therefore cannot function as a reason for me on which I base my belief. Yet I know I exist. My belief in my own existence, then, is arguably an instance of knowledge although it does not have a rational source. According to this objection, my belief has a source, but it is hidden from me. I'll discuss this objection next.

## 16. The problem of hidden sources

Sosa suggests that

the vast bulk of our present beliefs are not ones we hold for sufficient reasons motivationally operative at this very moment when they are held (mostly in storage). What is more, the vast bulk of these beliefs are ones whose rational aetiology is now beyond our ability to recall. So, for this vast bulk, it is hard to see how we could presently know their sources to be reliable, since we now lack access to what those sources are.<sup>30</sup>

The types of beliefs Sosa is talking about are stored beliefs. If Sosa is right that the etiology of stored beliefs is typically beyond our ken, then KR is a recipe for skepticism regarding stored beliefs.

In response, I suggest that we should distinguish between the distant causal basis operative when our beliefs were formed in the past, and the causal basis that sustains our beliefs in the present. The latter can be identified with seemings. Michael Huemer, for example, would say that my stored belief that water is H<sub>2</sub>O is causally sustained by a seeming: it seems to me that water is H<sub>2</sub>O.<sup>31</sup> I see no reason to think that reliability evidence for a seeming of this type is unavailable. The same reasoning applies to other stored beliefs. Typically, they are causally sustained by seemings we have each time a stored belief becomes an occurrent belief. And we have evidence that such seemings are by and large reliable. For example, many of these seemings are like the following: it seems to me that p because:

- I learned p at school;
- I heard it from an expert;
- I acquired the belief that p from a trustworthy source.<sup>32</sup>

To sum up, I agree with Sosa that stored beliefs pose a challenge for credentialist views that endorse the KR principle. At the same time, I believe there are resources of which advocates of such views can avail themselves to explain how we can know that stored beliefs have a reliable source.

## 17. Sosa's response to the circularity problem

The puzzle of reliability knowledge raises the twin problems of easy knowledge and potentially vicious circularity. If we reject dogmatism because we consider the consequence of easy knowledge unacceptable, we must accept some version of KR. Since it is difficult to see how perceptual reliability can be established a priori, we must accept some version of RK as well. And now we face the problem of source circularity: we are bound to use a source to confirm its own reliability. Fumerton thinks that's unacceptable:

You cannot use perception to justify the reliability of perception! You cannot use memory to justify the reliability of memory! You cannot use induction to justify the reliability of induction! Such attempts to respond to the skeptic's concerns involve blatant, indeed pathetic, circularity.<sup>33</sup>

<sup>30</sup>Sosa (2009: 217).

<sup>31</sup>This seeming is, like the belief, stored. Both rise to the level of consciousness when I consider the matter.

<sup>32</sup>See "Section 3: Conservatism and Lost Evidence" in Christensen (1994).

<sup>33</sup>Fumerton (1995: 177).

In the same vein, Sosa says:

Suppose one is disposed to acquire and sustain beliefs in a certain way, and suppose one wonders whether this disposition reliably yields beliefs that are true. It would seem viciously circular to answer that is indeed reliable ... based just on the deliverances of that very disposition.<sup>34</sup>

The general worry is clear, whether we are talking about seemings, faculties, competences, or dispositions: using a source to confirm its own reliability seems unacceptable. But why? Sosa suggests that, for a faculty to deliver epistemic goods, it must have proper epistemic status.<sup>35</sup> Sosa's point, I take it, is that a faculty cannot epistemically certify itself: it cannot confer proper epistemic status upon itself.

Let's apply this point to perception. To establish that perception is reliable, we must use perception. But if perception is to deliver epistemic goods – justified belief in its own reliability – we must already be in a position to justifiably presuppose that perception is reliable. How can we break out of this circle? Here is what I think is Sosa's response. First, he says that "it is far from obvious, and seems in fact false, that our epistemic competences generally operate through such presuppositions."<sup>36</sup> And then he moves on to say:

After all, if a competence might even do its work sub-personally, it is wildly implausible that the subject of that competence should be blocked (by some worry about vicious circularity) from discovering the specifics and even the reliability of the sub-personal mechanisms involved, from discovering this indeed through the deliverances of that very competence. Thus compare the discovery of how our vision works reliably and what makes it reliable, a discovery based essentially (not just causally but normatively) on the visual observations of scientists.<sup>37</sup>

Sosa's solution seems to be that a faculty or competence such as vision can, at a sub-personal level, deliver justified beliefs without a prior justified presupposition of reliability. *Prior* to having evidence of visual reliability, we can *already* use vision to generate justified beliefs that permit us to then infer that vision is reliable. So, in effect, Sosa is suggesting that, as far as vision is concerned, we must reject KR. But then it is unclear why vision isn't subject to the problem of easy knowledge in just the way Cohen says it is. And Sosa seems to agree with Cohen that the problem of easy knowledge is pernicious. Sosa, then, would have to explain how vision can be exempted from the KR principle without thereby allowing for the creation of easy knowledge.

### 18. Diachronic-successive versus synchronic-structural epistemic priority

I would like to suggest a different solution to the circularity problem. I think we should accept both KR and RK and reject

NP If KR and RK are both true, then neither perceptual knowledge nor knowledge of perceptual reliability are possible.

<sup>34</sup>Sosa (2009: 235).

<sup>35</sup>Sosa (2009: 236).

<sup>36</sup>Sosa (2009: 235).

<sup>37</sup>Sosa (2009: 236f).

KR and RK both assert priority relations. The following two articulations of KR and RK makes that explicit:

- KR<sub>pri</sub> Evidence of perceptual reliability is prior to perceptual knowledge: for a perceptual experience to be a knowledge source, we must *already* have prior evidence of perceptual reliability.
- RK<sub>pri</sub> Perceptual knowledge is prior to having perceptual evidence: for us to acquire perceptual evidence, perceptual experience must *already* be a knowledge source.

NP is plausible if we interpret the priority relations KR<sub>pri</sub> and RK<sub>pri</sub> assert as diachronic-successive priority: one thing must happen before the other thing can happen. But there is an alternative. I believe we should think of the priority relations in question as synchronic-structural: one thing cannot happen without the other thing being present *at the same time*. Understood this way, the two principles may be recast thus:

- KR<sub>sim</sub> You cannot have perceptual knowledge without having, at the same time, evidence of perceptual reliability.
- RK<sub>sim</sub> You cannot have evidence of perceptual reliability without perception being, at the same time, a knowledge source.

KR<sub>sim</sub> and RK<sub>sim</sub> do not say that you must have evidence of perceptual reliability *before* you can have perceptual knowledge, and that you must already have perceptual knowledge *before* you can acquire evidence of perceptual reliability. Rather, the principles say merely that you cannot have one without the other. This allows for the possibility that both conditions are satisfied at the same time, with neither one of them being the sort of condition that must be met before the other one can be met. The two conditions need not – and indeed cannot – be satisfied successively. They only need to be satisfied simultaneously.

I suggest that KR<sub>pri</sub> and RK<sub>pri</sub> should be construed as KR<sub>sim</sub> and RK<sub>sim</sub>. Next, I will argue that there is good reason to think that the circularity involved in mutual dependence of this kind is benign.

### 19. Benign mutual dependence: the clown fish and the sea anemone

The world is actually full of examples of benign mutual dependence. Consider the ecosystem in which we find the clown fish and the sea anemone. Both of these creatures depend on each other in various ways.<sup>38</sup> Here is an argument that, duplicating the structure of the puzzle of reliability knowledge, focuses on the mutual dependence between the clownfish and the sea anemone:

- FA The clownfish cannot exist without the prior existence of the sea anemone.
- AF The sea anemone cannot exist without the prior existence of the clown fish.
- NP If FA and AF are both true, then neither the clown fish nor the sea anemone can exist.

<sup>38</sup>For example, they protect each other against predators.

Therefore:

BN Neither one can exist.

But these creatures do exist. So what's wrong with the argument? Here is the solution: FA and AF must not be understood as asserting conditions that must be satisfied successively. Rather, they assert a relationship of structural mutual dependence: the two conditions must merely be met at the same time. If FA and AF are understood in this way, NP is false.

There are many other examples that illustrate the same point. The organs of an organism depend on each other. It's not the case that an organism's organs must come into play one after the other. Rather, for an organism to be alive, they must all be in place at the same time. The same goes for the parts of an internal combustion engine: cylinders, spark plugs, carburetor, etc. It's not the case that the cylinders must start working before the carburetor does, or that the carburetor must already be doing its job before the pistons in the cylinders start moving. Rather, for the engine to be working, all its parts must be in functioning order at the same time.

I want to suggest that we should think of perceptual justification and evidence of perceptual reliability as necessary parts of the cognitive engine. For the engine to produce perceptual knowledge, they must both be in place at the same time. But neither of them needs to precede the other one. Mutual dependence of this kind involves source circularity: perception confirms the reliability of perception, and evidence of perceptual reliability certifies perception as a justification source. That's circularity all right, but there is nothing vicious about it, just as there is nothing vicious in the way the clown fish and the sea anemone, a heart and a liver, or a cylinders and spark plugs depend on each other.<sup>39</sup>

## 20. Easy knowledge and Sosa's two-stage model of reflective knowledge

While Sosa is known for advocating that there is such a thing as animal knowledge, in the end, it seems to me, what he really cares about is reflective knowledge. Reliably functioning faculties give us animal knowledge. Reflective knowledge involves more than reliably functioning faculties: it requires knowledge that our faculties are reliable. Thus Sosa faces the puzzle of reliability knowledge and its twin problems: easy knowledge and seemingly vicious circularity. What is his solution? It is, it would seem, a two-stage model that is diachronic-successive. At the first stage, vision delivers epistemic goods although we do not yet have justification for taking it to be reliable. At the second stage, enough epistemic goods are in place to recognize the reliability of vision. Visual knowledge then becomes reflective knowledge.

Unlike Pryor, Sosa does not argue that easy knowledge is benign.<sup>40</sup> But the two-stage model does not steer clear of easy knowledge. If vision can give me epistemic goods although I do not yet have evidence of its reliability, then I can use these epistemic goods for dismissing deceptive alternatives. I share Cohen's intuition that that's epistemically defective reasoning, and it looks like Sosa shares this intuition too. It would seem, therefore, Sosa needs to explain how the two stage model he advocates avoids the problem of easy knowledge.

<sup>39</sup>The argument I have offered is an argument from analogy. Critics might argue that the analogy fails because there is a relevant difference. In ecosystems, organisms, and engines, the kind of mutual dependence in question is purely factual. But the mutual dependence of perceptual knowledge and knowledge of perceptual reliability is normative. My response is that, while that is of course an interesting difference, it is unclear why it should establish that the mutual dependence in the latter case is not benign. For the objection to stick, details would have to be offered.

<sup>40</sup>See Sosa (2009: note 23 on p. 230). There, Sosa rejects Pryor's (2004) attempt to defuse the charge of easy knowledge.

## 21. The web

What makes knowledge reflective is having justification for believing in, and ideally knowing, the reliability of our sources. For externalists, reliability knowledge is seemingly unproblematic: if perception produces reliably the true belief that perception is reliable, you know that perception is reliable. For dogmatists, too, reliability knowledge is seemingly unproblematic. Perception gives us knowledge without antecedent evidence of its reliability. Using this knowledge, we can then bootstrap to knowledge of perceptual reliability. According to the objection from easy knowledge, each of these views makes knowledge or perceptual reliability too easy to acquire. The alternative is credentialism: perception can give us knowledge only if we have, at the same time, evidence of its reliability. On this approach, all perceptual knowledge is reflective – or at least quasi-reflective if reflective knowledge requires actual reflection. While credentialism avoids the problem of easy knowledge, it runs into another problem: that of source circularity, or self-support. Sosa endorses the metaphor of the self-supporting web as the right way of understanding how reflective knowledge comes about. He says:

The right model for understanding reflective justification is not the linear model whereby justification is a sort of liquid that flows through some pipe or channel of reasoning, from premises to conclusion ... A better model is rather that of the web of belief, whereby the web is properly attached to the environment, while its nodes can also gain status through mutual support.<sup>41</sup>

But here is a question about the epistemic role the web plays. Does the web *create* epistemic justification, or does it merely *augment* previously created justification? It seems to me that, unless there is a way of defusing the charge of easy knowledge, we must endorse KR. But if KR holds without exception, then there is no perceptual knowledge at all without evidence of perceptual reliability. If we take that to be true, we must conclude that there is no perceptual justification outside of the web. The web creates all the perceptual justification there is.

## 22. The creation of epistemic justification

Can the web create epistemic justification all by itself? One reason for saying ‘no’ is that the web can’t by itself deliver the initial boost. Sosa suggests a second reason. According to him, the answer is ‘no’ because hierarchically basic knowledge is unavoidable. It is unavoidable because, he says,

it is impossible for humans actually to ascend infinite ladders of reflection. Reflection is bound to give out at some level, where we no longer have an explicit or implicit belief of the form  $B^{n+1}p$  although we do have a belief of the form  $B^n p$  (where the superscripts represent the number of levels involved).<sup>42</sup>

This argument is sound if we take the ‘know’ in the consequent of the knowledge version of KR, namely

$KR_K$  A potential knowledge source, S, can yield K for you → you know that S is reliable.

<sup>41</sup>Sosa (2009: 239).

<sup>42</sup>Sosa (2009: 228).

to refer to explicit knowledge requiring occurrent belief. As I mentioned early on in this paper, if we interpret KR in this way, KR is not plausible. But from the implausibility of KR thus understood, it doesn't follow that there aren't less demanding and thus plausible versions of the principle. I believe that, if we interpret the knowledge operator in the consequent to mean "you must have knowledge-grade evidence for the reliability of S," then KR is plausible. On this approach, KR does not generate an infinite level of reflection. Therefore, I don't agree with Sosa's claim that hierarchically basic knowledge is unavoidable.

Nor do I think that the initial boost objection is compelling. If we adopt the synchronic-structural understanding of epistemic priority, then there is no need for an initial boost. During the cognitive development of a child, perceptual knowledge and knowledge of perceptual reliability can emerge together just as, during the evolution of sea life in shallow lagoons and sheltered reefs, the clown fish and the sea anemone evolved together. On this model, we have gradual emergence instead of an initial boost.

How, then, does perceptual justification gradually emerge? When an infant starts observing the world around her, she notices that what can be seen can be touched, heard, tasted, and smelled. What she saw yesterday she sees again today. A rattle that was red yesterday is still red today. With enough effort, one can fill pages describing similar ways in which perceptual submodules confirm each other. As the infant develops, her cognitive web grows as well and includes introspection, reason, and memory. The latter in particular plays an essential role since, to a large extent, evidence of the reliability of a source is supplied by remembering its performance in the past.

What I have described is of course just an initial and rather rough account of the basic idea: the web creates perceptual justification through a continuous flux of mutual confirmations, which, by virtue of being remembered, constitute evidence of cognitive reliability.<sup>43</sup> At the beginning of cognitive development, as soon as the infant starts trusting her perceptual powers she remembers perceptual success in the past, and perceptual justification gradually emerges. An essential element of this picture is that belief sources can confirm their own reliability. This is particularly striking in the case of memory, which, through the remembered track record of its past performance, confirms its own reliability. Like Sosa, I see no reason to think that there is anything wrong with epistemic self-support, thus conceived. Sosa, however, seems to think that epistemic self-support merely augments perceptual justification that was previously created through hierarchical basicness. Unlike Sosa, I believe the lesson of the problem of easy knowledge is that there is no hierarchically basic knowledge. The web creates all the perceptual justification there is.<sup>44</sup>

## References

- BonJour L.** (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Brogaard B.** (2013). 'Phenomenal Seemings and Sensible Dogmatism.' In C. Tucker (ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, pp. 270–89. Oxford: Oxford University Press.

<sup>43</sup>For a more detailed account, see Steup (2019).

<sup>44</sup>A previous version of this paper was presented at the 14th *Episteme* Conference in honor of Ernest Sosa, July 24–25, 2018, in Tenerife, Canary Islands. Thanks to Jennifer Lackey for organizing this terrific conference. For helpful discussion, feedback, and in some cases correspondence, I'm indebted to Laura Callahan, Jonathan Ichikawa, Thomas Grundman, Peter Klein, Ram Neta, Baron Reed, Lisa Miracchi, Joshua Schechter, and Ernest Sosa.

- Brown J.** (2005). 'Doubt, Circularity and the Moorean Response to the Sceptic.' *Philosophical Perspectives* 19, 1–14.
- Chisholm R.** (1980). *The Foundations of Knowing*. Minneapolis, MN: University of Minnesota Press.
- Christensen D.** (1994). 'Conservatism in Epistemology.' *Noûs* 28, 69–89.
- Cohen S.** (2002). 'Basic Knowledge and the Problem of Easy Knowledge.' *Philosophy and Phenomenological Research* 65, 309–29.
- Cohen S.** (2005). 'Why Basic Knowledge is Easy Knowledge.' *Philosophy and Phenomenological Research* 70, 417–30.
- Cohen S.** (2010). 'Bootstrapping, Defeasible Reasoning, and *A Priori* Justification.' *Philosophical Perspective* 24, 142–59.
- Davies M.** (2004). 'Epistemic Entitlement, Warrant Transmission and Easy Knowledge.' *Aristotelian Society* Suppl. Vol. 78, 213–45.
- Feldman R.** (2003). *Epistemology*. Upper Saddle River, NJ: Pearson.
- Fricker E.** (1994). 'Against Gullibility.' In B.K. Matilal and A. Chakrabarti (eds), *Knowing From Words*, pp. 125–61. Dordrecht: Kluwer.
- Fumerton R.** (1995). *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Harman G.** (1986). *Change in View*. Cambridge, MA: MIT Press.
- Huemer M.** (2001). *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.
- Lackey J.** (2008). *Learning from Words*. New York, NY: Oxford University Press.
- Lackey J. and Sosa E.** (2006). *The Epistemology of Testimony*. Oxford: Oxford University Press.
- Lycan W.G.** (1988). *Judgement and Justification*. Cambridge: Cambridge University Press.
- Lycan W.G.** (2013). 'Phenomenal Conservatism and the Principle of Credulity.' In C. Tucker (ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, pp. 293–305. Oxford: Oxford University Press.
- Markie P.** (2005). 'Easy Knowledge.' *Philosophy and Phenomenological Research* 70, 406–16.
- Neta R.** (2010). 'Liberalism and Conservatism in the Epistemology of Perceptual Belief.' *Australasian Journal of Philosophy* 88, 685–705.
- Pryor J.** (2000). 'The Sceptic and the Dogmatist.' *Noûs* 34, 517–49.
- Pryor J.** (2004). 'What's Wrong with Moore's Argument?' *Philosophical Issues* 14, 349–78.
- Quine W.V.O.** (1951). 'Two Dogmas of Empiricism.' *Philosophical Review* 60, 20–43.
- Silins N.** (2007). 'Basic Justification and the Moorean Response to the Sceptic.' In T. Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology*, Vol. 2, pp. 108–42. Oxford: Oxford University Press.
- Sosa E.** (2009). *Reflective Knowledge*. Oxford: Oxford University Press.
- Steup M.** (2004). 'Internalist Reliabilism.' *Philosophical Issues* 14, 404–425.
- Steup M.** (2018). 'Destructive Defeat and Justificational Force: The Dialectic of Dogmatism, Conservatism, and Meta-Evidentialism.' *Synthese* 197, 2907–33.
- Steup M.** (2019). 'Benign Infinity.' In B. Fitelson, R. Borges and C. Braden (eds), *Themes from Klein: Knowledge, Scepticism, and Justification*, pp. 235–257. Cham: Springer.
- Tucker C.** (2010). 'Why Open-Minded People Should Endorse Dogmatism.' *Philosophical Perspectives* 24, 529–45.
- Vahid H.** (2004). 'Varieties of Epistemic Conservatism.' *Synthese* 141, 97–122.
- Vogel J.** (2000). 'Reliabilism Leveled.' *Journal of Philosophy* 97, 602–23.
- White R.** (2006). 'Problems for Dogmatism.' *Philosophical Studies* 131, 525–57.
- Wright C.** (2004). 'On Epistemic Entitlement: Warrant for Nothing (and Foundations for Free)?' *Aristotelian Society Supplement* 78, 167–212.

**Matthias Steup** received his PhD in 1985 at Brown University. He has taught at Grinnell College, the University of Wyoming, St. Cloud State University, and Purdue University, and is since 2016 Chair of the Philosophy Department at the University of Colorado, Boulder. He is the author of *An Introduction to Contemporary Epistemology* (Prentice Hall, 1995) and numerous articles on knowledge and justification. He is a co-editor (with Ernest Sosa and Jonathan Dancy) of the *Blackwell Companion to Epistemology*, and a co-editor (with Ernest Sosa and John Turri) of *Contemporary Debates in Epistemology* (Wiley-Blackwell).