

SOLVING THE PROBLEM OF EASY KNOWLEDGE

TIM BLACK

In *The Philosophical Quarterly* 58 (2008): 597-617

(This version might be different in certain respects from the published version.)

Stewart Cohen argues that several epistemological theories fall victim to the problem of easy knowledge—they allow us far too easily to know that certain skeptical hypotheses are false and that the way things seem is a reliable indicator of the way they are. This problem is a result of the theories' interaction with an epistemic closure principle. Cohen chooses to focus on the theories, suggesting that we modify them in certain ways in order to avoid the problem of easy knowledge. I argue, on the other hand, that attempts to solve the problem should focus on closure and, in particular, that we can use a new and plausible epistemic closure principle in solving the problem of easy knowledge. My solution allows us to provide a uniform response to different versions of the problem of easy knowledge, and it succeeds where other proposed solutions fail.

Stewart Cohen argues that several epistemological theories fall victim to what he calls the problem of easy knowledge.¹ He maintains that this problem arises for theories that reject

(KR) A potential knowledge source K can yield knowledge for S only if S knows that K is reliable,

and thus for theories that allow for what he calls basic knowledge, which is knowledge that S acquires from K even when she fails to know that K is reliable. Cohen argues that on theories that allow for basic knowledge—that is, theories with a basic knowledge structure, or *BKS theories*—bootstrapping and closure allow us far too easily to know that certain skeptical hypotheses are false and that the way things *seem* is a reliable indicator of the way they *are*. This is the problem of easy knowledge.

As a solution to this problem, Cohen follows Ernest Sosa in distinguishing between animal knowledge and reflective knowledge.² Cohen then suggests that animal knowledge is not

¹ See Cohen, 'Basic Knowledge and the Problem of Easy Knowledge', hereafter BKPEK, *Philosophy and Phenomenological Research*, 65 (2002), pp. 309-329.

closed under known logical entailment, which means that closure won't be able to take us too easily from animal knowledge to, say, the knowledge that certain skeptical hypotheses are false. As a solution to the bootstrapping version of the problem, Cohen suggests that animal knowledge cannot combine individually with self-knowledge in order to provide the evidence necessary for knowing, for example, that the way things seem is a reliable indicator of the way they are. This keeps bootstrapping from giving us knowledge too easily.

Moreover, Peter J. Markie argues that even those theories that accept (KR)—that is, theories that do not allow for basic knowledge, or *non-BKS theories*—are susceptible to the problem of easy knowledge.³ There are cases in which I know that K is reliable and in which I know via K that some proposition *p* is true, but in which closure allows me too easily to know that some skeptical hypothesis—one that is incompatible with *p*—is false. In proposing a solution to this problem, Cohen appeals to epistemological contextualism.⁴ As we will see in Section 2, contextualism solves the problem by maintaining that there are no contexts in which we both know and know too easily—we fail to know at all in contexts in which the epistemic standards are high, and we know, but not too easily, in contexts in which the epistemic standards are more relaxed.

I propose an alternative solution to the problem of easy knowledge, an alternative with a focus that is different from that of Cohen's proposed solutions—and, in fact, different from the

² See BKPEK, pp. 326-327. For Sosa on the difference between animal and reflective knowledge, see, for example, his 'Knowledge and Intellectual Virtue', in E. Sosa, *Knowledge in Perspective: Selected Essays in Epistemology* (Cambridge UP, 1991), pp. 225-244; and E. Sosa, 'Two False Dichotomies: Foundationalism/Coherentism and Internalism/Externalism', in Walter Sinnott-Armstrong (ed), *Pyrrhonian Skepticism* (Oxford UP, 2004), pp. 146-160.

³ See Markie, 'Easy Knowledge', hereafter EK, *Philosophy and Phenomenological Research*, 70 (2005), pp. 406-416.

⁴ See 'Why Basic Knowledge is Easy Knowledge', hereafter WBKEK, *Philosophy and Phenomenological Research*, 70 (2005), pp. 417-429. For statements of his contextualism, see S. Cohen, 'How to be a Fallibilist', *Philosophical Perspectives*, 2 (1988), pp. 91-123; and S. Cohen, 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives*, 13 (1999), pp. 57-89.

focus of any solution yet proposed.⁵ As Cohen suggests, the problem of easy knowledge results from closure's interaction with theories of a certain sort, and this holds even for the bootstrapping and non-BKS versions of the problem. On the bootstrapping version, closure maintains that since I know both that the table is red and that it looks red, I know that its looking red accurately indicates its being red. And bootstrapping brings knowledge that one thing *reliably* indicates another only if I've acquired in a sufficient number of particular cases the knowledge that one thing *accurately* indicates another. On the non-BKS version, closure maintains that since I know that the table is red, I know that some skeptical hypothesis is false. Closure is essential, then, to each version of the problem of easy knowledge. Even so, Cohen's solutions focus on the *theories*—he argues that BKS theories should distinguish between animal knowledge and reflective knowledge, and that non-BKS theories should adopt contextualism. My solution, on the other hand, focuses on closure, and as a consequence, it requires us neither to revise nor to supplement any theory: On my solution, we need neither to adopt contextualism nor to distinguish between animal and reflective knowledge. I argue that we can use a particular epistemic closure principle in solving the problem of easy knowledge as it arises both for BKS theories and for non-BKS theories. Unlike Cohen's solution, then, mine allows us to provide a uniform response to different versions of the problem of easy knowledge. This is to be preferred since the problem, even as it arises in apparently different ways, springs from a single source—namely, closure—and my solution focuses on that very source.⁶ Moreover, since, as I argue in

⁵ See footnotes 6 and 7.

⁶ Davies, who argues that easy-knowledge inferences are not sound, proposes a solution based on his Limitation Principle (see M. Davies, 'A Principled Solution to the Problem of Armchair Knowledge', Third Hempel Lecture, Princeton University, Friday 26 September 2003 [online at <http://philrssi.anu.edu.au/~mdavies/papers/Hempel3.pdf>], pp. 2, 12-13), which he offers as a supplement to epistemological theories. Although I agree with Davies—and with Cohen—that easy-knowledge inferences are not sound, I suggest that we focus on closure rather than on theories.

Section 1, Cohen's proposed solution fails to solve the BKS version of the problem, my solution succeeds where his does not.⁷

1. THE PROBLEM OF EASY KNOWLEDGE FOR BKS THEORIES

There are two versions of the problem of easy knowledge for BKS theories, one that involves closure, and another that involves bootstrapping. We can illustrate the latter in terms of evidentialist foundationalism, and the former in terms of reliabilism.

1a. Closure

According to a reliabilism that rejects (KR), 'I do not need to know my perceptual processes are reliable in order for me to look at [a] table and thereby come to know it is red. It's enough that my perceptual processes are reliable' (BKPEK, p. 315). Now consider an epistemic closure principle, which I'll call *traditional closure*, according to which 'If S knows P and S knows P entails Q, then S knows (or at least is in a position to know) Q' (BKPEK, p. 312). Given traditional closure, the knowledge that certain skeptical hypotheses are false seems to come all too easily with, for example, our basic knowledge that the table is red. Suppose that I know both that the table is red, and that if it's red, then it isn't white and illuminated by red lights. It

⁷ Other strategies for solving the problem of easy knowledge include arguing that easy-knowledge inferences are sound even though they seem unsound (see P. Klein, 'Closure Matters: Academic Skepticism and Easy Knowledge', *Philosophical Issues*, 14 (2004), pp. 165-184; J. Pryor, 'What's Wrong with Moore's Argument?', *Philosophical Issues*, 14 (2004), pp. 349-378; and Markie in EK) and arguing that easy-knowledge inferences are sound in some contexts but unsound in others (see M. Bergmann, 'Epistemic Circularity: Malignant and Benign', *Philosophy and Phenomenological Research*, 69 (2004), pp. 709-727; and R. Neta, 'A Contextualist Solution to the Problem of Easy Knowledge', *Grazer Philosophische Studien*, 69 (2005), pp. 183-205). The former strategy, as Cohen suggests, doesn't rid us of the intuition that something is wrong in cases in which we seem to know too easily (see WBKEK, pp. 418-420). To those who adopt the latter strategy, we can say, as Cohen does, that the problem of easy knowledge arises even in contexts in which easy-knowledge inferences are supposed to be unsound (see BKPEK, pp. 314-315, n. 15; and WBKEK, pp. 424-427).

follows by traditional closure that I know that the table isn't white and illuminated by red lights. It seems, however, that this knowledge comes much too easily.⁸

In proposing a solution to this problem, Cohen follows Sosa in distinguishing animal knowledge from reflective knowledge. According to Sosa, '[o]ne has *animal knowledge* about one's environment, one's past, and one's own experience if one's judgments and beliefs about these are direct responses to their impact – e.g., through perception or memory – with little or no benefit of reflection or understanding', while '[o]ne has *reflective knowledge* if one's judgment or belief manifests not only such direct response to the fact known but also understanding of its place in a wider whole that includes one's belief and knowledge of it and how these come about' ('Knowledge and Intellectual Virtue', p. 240).

After endorsing this distinction, Cohen suggests that animal knowledge and reflective knowledge are different in kind because they play different roles in reasoning. For one thing, we can suppose that animal knowledge is not closed under known entailment, but that reflective knowledge is (see BKPEK, p. 327). Given this, our having animal knowledge that the table is red does not entail, as traditional closure suggests it does, that we know that the table isn't white and illuminated by red lights. This keeps us from knowing too easily via traditional closure that certain skeptical hypotheses are false.

But Cohen's solution is flawed. To see this, note first that when we argue that animal knowledge isn't closed under known entailment, we can focus either on skeptical cases or on non-skeptical cases. Yet neither sort of case supports Cohen's solution. Consider non-skeptical

⁸ See BKPEK, p. 313. Keep in mind that closure principles are different from transmission principles. According to transmission, warrant transmits from S's standing with respect to *p* and her standing with respect to *p* entails *q* to her standing with respect to *q*, and it does so in such a way that by deducing *q* from *p* and *p* entails *q*, S can come to know that *q* for the first time. Closure is weaker than transmission. It simply provides a necessary condition on the truth of a certain conjunction, maintaining that S knows both that *p* and that *p* entails *q* only if S knows that *q*. I am concerned *only with closure*; no principle proposed here requires epistemic agents to engage in an act of deduction.

cases, which include cases like one in which I know both that the table is red and that if it's red, then it will match the carpets in my dining room. Given these two bits of knowledge, BKS theories maintain that I also know that the table will match the carpets in my dining room. This application of traditional closure seems entirely acceptable. If our thinking about closure and animal knowledge is guided only by cases like this, we have no reason to conclude that animal knowledge isn't closed under known entailment.

In order to support Cohen's solution, then, we must focus on skeptical cases, such as the one he in fact considers. Such cases support the claim that animal knowledge isn't closed: Even when I know both that the table is red and that if it's red, then it isn't white and illuminated by red lights, it nonetheless seems that I *don't* know that the table isn't white and illuminated by red lights. In fact, as Markie notes, even if I have *reflective* knowledge that the table is red—that is, even if I know that the source of that knowledge is reliable—it still seems that I don't know that the table isn't white and illuminated by red lights (see EK, p.410). So, if we let skeptical cases guide our thinking, it seems that we should conclude that neither animal knowledge *nor reflective knowledge* is closed under known entailment. According to Cohen's proposal, though, reflective knowledge is different from animal knowledge in just this respect. Thus, skeptical cases, just like non-skeptical cases, provide no support for Cohen's solution to the problem of easy knowledge.

1b. Bootstrapping

To illustrate how bootstrapping gives rise to the problem of easy knowledge for BKS theories, consider evidentialist foundationalism. According to evidentialist foundationalisms that reject (KR),

one can know P on the basis of evidence E without knowing that E is a reliable indication of P. For example, one can know that X is red, on the basis of its looking red, without knowing that X's looking red is a reliable indication of X's being red. (BKPEK, pp. 309-310)

Suppose, then, that I know both that the table is red and that it looks red. 'So I now have *some* evidence that something's looking red is a reliable indication that it is red. And by taking a few more looks I can acquire more evidence'(BKPEK, p. 318). Eventually, I will have acquired enough evidence to bootstrap up to the knowledge that something's looking red is a reliable indication of its being red. Yet it seems that this knowledge comes too easily.

In proposing a solution to this problem, Cohen once again appeals to the distinction between animal knowledge and reflective knowledge. In this case, animal knowledge and reflective knowledge are supposed to play different roles in reasoning because 'animal knowledge cannot combine individually (non-holistically) with self-knowledge to generate inferences' (BKPEK, p. 327). Given this, bootstrapping, as it is described above, cannot supply the knowledge that something's looking red is a reliable indication of its being red, for my animal knowledge that the table is red cannot combine individually with my self-knowledge that the table looks red in order to generate the inferences that are required for bootstrapping.

Once again, however, this solution is flawed, for there are cases in which animal knowledge *does* combine individually with self-knowledge to generate inferences.⁹ Suppose that I know via vision that I've just cut my finger. Suppose too that I know via the somatosensory modality that it now feels to me as if I've just cut my finger. Intuitively, it seems perfectly reasonable to suppose in this case that these two pieces of knowledge support a third, *viz.*, that its feeling to me as if I've just cut my finger accurately indicates that I've just cut my finger. (I use

⁹ For a similar criticism, see EK, p. 413.

‘accurate’ and its cognates in discussing *particular instances* in which something’s seeming to have a certain feature indicates its having that feature. Given this, if something’s seeming to have a certain feature *reliably* indicates its having that feature, then its seeming to have that feature must *accurately* indicate, in at least the majority of cases, that it has that feature.)

Let me be clear about the details of this case. In

(1) I know via vision that I’ve just cut my finger,

the propositional content of my knowledge state, which is delivered by vision, is just this: *I’ve just cut my finger*. It seems completely appropriate to characterize this knowledge as a piece of animal knowledge, for my belief that I’ve just cut my finger is a direct response, with little or no help from reflection or understanding, to the impact of my experience.

Furthermore, we may think of

(2) I know via the somatosensory modality that it now feels to me as if
I’ve just cut my finger

as a piece of self-knowledge. The propositional content of this knowledge state, which is delivered by my somatosensory faculty, is just this: *it now feels to me as if I’ve just cut my finger*. This is precisely the sort of self-knowledge with which Cohen is concerned, for it represents my knowledge of how something perceptually seems to me to be.

Finally, given both (1) and (2), it seems that

(3) I know that its feeling to me as if I’ve just cut my finger accurately
indicates that I’ve just cut my finger.

This fits the pattern of Cohen’s case: (1) and (2) support a piece of knowledge that concerns whether something’s seeming to have a certain feature accurately indicates its having that feature, and this is just the sort of knowledge that is required for bootstrapping. Here, then, is a

case in which the animal knowledge in (1) combines individually with the self-knowledge in (2) to support my knowing that something's seeming to have a certain feature accurately indicates its having that feature.

But maybe this is all too quick. Suppose that a cautious individual confronts me. Suppose, for example, that my daughter wants to know why I'm dancing around holding my finger and grimacing wildly.¹⁰ I say to her, 'I've just cut my finger.' She's a careful one, though, and so she worries, 'Daddy, what if someone is just making it *feel* as if you've cut your finger—with, like, some weird drug or something—when you didn't *really* cut your finger?' I reply, 'Oh, don't worry. I *see* that I've cut my finger—I saw the knife's blade move across my finger, which coincided roughly with the onset of the pain, and now I see blood issuing from a gash in my finger—so no one is just making it feel as if I've cut my finger.' I think that, in this case, my daughter will be satisfied with my response. After all, she can herself see blood issuing from the gash in my finger.

Even if my daughter confronts me about the accuracy of vision, a similar response would satisfy her. She might say, 'What if someone is just making it *look* as if you've cut your finger, maybe with some carefully applied ketchup?' I reply, 'Don't worry. The liquid on my finger neither smells nor tastes like ketchup—in fact, it smells and tastes like blood—and the ketchup has been in the refrigerator all this time.' I think my daughter would be satisfied with this response, and that she would allow that I know in this case that its looking to me as if I've just cut my finger accurately indicates that I've just cut my finger.

I conclude that there are legitimate counterexamples to Cohen's solution to the bootstrapping version of the problem of easy knowledge. We've also seen that there is no

¹⁰ The examples in this paragraph and in the next are meant to remind us of—and to stand in contrast to—those presented in BKPEK, pp. 314-315.

support to be found for his solution to the closure version of the problem. We must therefore look elsewhere for a solution to the problem of easy knowledge as it arises for BKS theories.

2. THE PROBLEM OF EASY KNOWLEDGE FOR NON-BKS THEORIES

The problem of easy knowledge is also more general—it applies to non-BKS theories as well as to BKS theories. Suppose that S has enough inductive evidence to know that her vision is reliable. According to non-BKS theories, S can then know on the basis of her vision that the table is red. Indeed, since S knows that her vision is reliable, her knowledge that the table is red counts as reflective knowledge, which is supposed to be closed under known entailment. Given, then, that S knows that the table's being red entails that it isn't white and illuminated by red lights, she knows that the table isn't white and illuminated by red lights. Here again, though, it seems that S comes by this knowledge too easily.

In solving this version of the problem, Cohen appeals to epistemological contextualism, according to which the truth conditions of knowledge attributions depend on certain features of the attributor's context. Contextualists contend that these features—and hence the truth conditions of knowledge attributions—can shift from context to context. Those truth conditions can be quite demanding in some contexts, so that we cannot truthfully attribute knowledge when we find ourselves in contexts like that. Yet the truth conditions for knowledge attributions are comparatively relaxed in other contexts, in which, as a consequence, we *can* truthfully attribute knowledge.

Contextualism can solve the non-BKS version of the problem of easy knowledge. In relaxed-standards contexts, S can have enough inductive evidence to know that her vision is reliable, she can know via vision that the table is red, and she can know that the table isn't white

and illuminated by red lights. This knowledge doesn't count as easy knowledge, though, because we expect that S will have that knowledge in relaxed-standards contexts. Yet in contexts in which the truth conditions for knowledge attributions are very demanding, S knows neither that the table is red nor that it isn't white and illuminated by red lights—and this is true even if she knows that her vision is reliable. Contextualism therefore solves the non-BKS version of the problem of easy knowledge by maintaining that the knowledge we acquire in relaxed-standards contexts is not acquired too easily, and that in high-standards contexts, we fail to know *at all* that certain things are true.

3. FOCUSING ON CLOSURE

Rather than adopting contextualism as a solution, however, I propose a different solution, one that, unlike contextualism, solves both the BKS and the non-BKS versions of the problem of easy knowledge. Now, any solution to this problem should answer the question, 'Why is it that, in the relevant range of cases, S cannot gain ... knowledge so easily?' (Neta, 'A Contextualist Solution to the Problem of Easy Knowledge', p. 189). My answer goes something like this: In the relevant range of cases, the connections that are alleged to provide easy knowledge are connections between a bit of knowledge gained via one belief-producing mechanism (or set of belief-producing mechanisms) and a bit of knowledge gained at least in part via a *different* belief-producing mechanism (or set of belief-producing mechanisms).¹¹ Traditional closure tolerates such connections, and so rejecting traditional closure, as Cohen does, at least for some

¹¹ There are, of course, familiar difficulties with individuating belief-producing mechanisms. Here, however, I rely on the somewhat standard classification of mechanisms—perception (including vision, audition, taste, olfaction and touch (or somatosensation), along with proprioception, kinesthesia and the like), memory, testimony, reason (or intuition) and other less highly esteemed mechanisms such as clairvoyance. For a nice treatment of the individuation of sensory modalities, see B. Keeley, 'Making Sense of the Senses: Individuating Modalities in Humans and Other Animals,' *The Journal of Philosophy*, 99 (2002), pp. 5-28.

kinds of knowledge, affords a solution to the problem. Yet since we have the strong intuition that *some* epistemic closure principle is true—for one thing, some such principle might stand behind or certify our deductive practices—this solution to the problem of easy knowledge leaves us in a somewhat awkward position. It would be nice, then, to find a closure principle that licenses the connections that intuitively yield knowledge but that refuses to allow those that yield easy knowledge. One principle that does this, as I argue below, is a principle that licenses only connections between knowledge gained via one belief-producing mechanism (or set of such mechanisms) and knowledge gained via the *same* belief-producing mechanism (or set of such mechanisms). (To avoid the awkward parenthetical expressions employed in this paragraph, I will from now on, when I speak of one belief-producing mechanism, understand this to mean either one such mechanism or one set of such mechanisms.)

To establish the plausibility of such a principle, consider first the white-table hypothesis—the hypothesis that the table is white and illuminated by red lights. In the context of Cohen’s examples, in which the hypothesis calls into question the claim that the table’s looking red accurately indicates its being red, we should conceive of it as a skeptical hypothesis. This means, perhaps among other things, that it describes a scenario that is *ex hypothesi* visually indistinguishable from the actual scenario, and that its obtaining is incompatible either with the obtaining of the actual scenario or with one’s knowing that the actual scenario obtains. Now, to know that the white-table hypothesis is false, we must be able to distinguish the actual scenario, which I take to be one in which the table is red, from scenarios in which the white-table hypothesis is true. But such scenarios are *ex hypothesi* visually indistinguishable from the actual scenario. This means, for one thing, that I cannot know that the white-table hypothesis is false *simply by looking around and seeing* that the table is in fact red and illuminated only by white

lights—it is extremely difficult, to say the least, to know via vision that the white-table hypothesis is false.¹²

We can now make our way to a closure principle that will help us in solving the problem of easy knowledge. We start with the Independence Principle,¹³ which represents an attractive way of responding to cases in which we seem to have easy knowledge of the reliability of certain belief-producing mechanisms. We then continue through a series of corrections, adjustments and deductions, until we arrive finally at our closure principle. The Independence Principle says:

(IP) Where a belief gains its *prima facie* justification for S just from the fact that it was produced by a particular faculty (given, e.g. the faculty's reliability, proper function, ability to provide the subject with evidencing experiences), the belief is not supporting evidence for S for beliefs concerning the reliability of that very faculty. (EK, p. 414)¹⁴

Suppose that my belief that the table is red gains its *prima facie* justification just from the fact that it was produced by vision. Given (IP), this means that my belief that the table is red isn't supporting evidence for me for beliefs concerning the reliability of vision. But note that my belief that the white-table hypothesis is false is a belief that concerns the reliability of vision. If the white-table hypothesis were true, it would be the case both that the white table *looks* red to me and that I believe via vision that it *is* red. In that case, though, the table's looking red would not accurately indicate its being red, and this counts against the claim that something's looking red *reliably* indicates its being red. The white-table hypothesis therefore concerns the reliability

¹² Technically, of course, to know that the white-table hypothesis is false is to know that the disjunction, *either the table is not white or it is not illuminated by red lights*, is true. And this disjunction is true in plenty of scenarios in which the table is not red. Nevertheless, given that the white-table hypothesis is aimed directly at the claim that the table's looking red accurately indicates its being red, and given that the table does in fact look red, I take it that the only cases that matter in which the disjunction is true are those in which the table is red.

¹³ The Independence Principle is Markie's; see EK, p. 414.

¹⁴ Markie does not endorse a solution based on (IP).

of vision. Given this, (IP) provides no bridge from S's knowing via vision that the table is red to her knowing that the white-table hypothesis is false.

Markie suggests, however, that (IP) cannot help us to solve the problem of easy knowledge, for it runs afoul of the following 'intuitively attractive principle' (EK, p. 415):

(IAP) If it is reasonable for us to believe p , and if p counts in favor of q , then p is a reason (perhaps defeasible) for us to believe q .¹⁵

For suppose that vision makes it reasonable for me to believe that I'm seeing a red table, where the fact that I'm seeing a red table counts in favor of vision's accuracy, at least in conjunction with the fact that the table looks red. (IAP) then suggests, *contra* (IP), that my belief that I'm seeing a red table, which was produced by vision, is a reason for me to believe that vision is accurate in this case.

Yet even though Markie finds it intuitively attractive, (IAP) is false. Suppose that a particular car is red and that it is entirely reasonable for me to believe that it is red. Unbeknownst to me, however, Sarah, of whom I have no idea, strongly dislikes red cars. Its being red therefore counts in favor of her not buying the car. But its being red is not a reason for me to believe that Sarah won't buy the car, where to say that p is not a reason *for me* to believe that q is to say that I am not cognizant of the fact that p counts in favor of q . There is, of course, a weaker sense of 'reason' in which p is a reason to believe that q simply in virtue of the fact that it counts in favor of q . (In this weaker sense, the car's being red is a reason to believe that Sarah won't buy it.) But p 's counting in favor of q is far from sufficient for my being cognizant of its

¹⁵ See EK, p. 415. Markie's principle uses 'the truth of p increases the likelihood that another proposition, q , is true', where (IAP) uses ' p counts in favor of q '. I make this change, trusting that it takes nothing away from Markie's principle, because it is a mistake to put the principle in terms of an increase in likelihood. Suppose that in June, the probability that Candidate A will be elected is 0.1. Later, A changes her stance on an important issue so that the probability that she'll be elected increases to 0.2. The fact that A changed her stance on the issue increases the likelihood that she'll be elected, yet even if it's reasonable for me to believe that she's changed her stance, this doesn't constitute a reason for me to believe that she will in fact be elected. (Thanks here to Leo Iacono.)

doing so. Since (IAP) invokes reasons in the stronger sense, the counterexample above shows it to be false.

The counterexample to (IAP) is helpful in that it suggests an alternative principle:

(IAP*) If it is reasonable for S to believe that *p*, and if *it is reasonable for S to believe that p* counts in favor of *q*, then *p* is a reason (perhaps defeasible) for S to believe *q*.¹⁶

This principle accommodates the counterexample to (IAP). If it is reasonable for me to believe both that the car is red and that its being red counts in favor of Sarah's not buying it, then its being red is indeed a reason for me to believe that Sarah won't buy the car.

Given the success of (IAP*), we might try the following principle, which is formulated in terms of justification rather than in terms of reasons and reasonable belief.

(IAP*_J) If I justifiedly believe that *p* and that *p* counts in favor of *q*, then my belief that *p* justifies my belief that *q*.

From (IAP*_J), we can derive via substitution a principle that is preferable to (IP), a principle that better expresses the relationship between a belief's being justified via K and its providing support for beliefs about K's reliability:

(IP*) If I justifiedly believe via K that *p*, and if I justifiedly believe that *p* counts in favor of K's reliability, then my belief that *p* justifies my belief that K is reliable.

But what good does (IP*) do us? It's formulated in terms of justification, while I have said that my solution to the problem of easy knowledge will focus on closure, which is understood in terms of knowledge. Perhaps we can simply reformulate (IP*):

(IP*_K) If I *know* via K that *p*, and if I *know* that *p* counts in favor of K's reliability, then

¹⁶ Compare Sosa, 'Two False Dichotomies: Foundationalism/Coherentism and Internalism/Externalism', p. 147; and page 488 of R. Feldman, 'In Defence of Closure', *The Philosophical Quarterly*, 45 (1995), pp. 487-494.

my belief that p supports my *knowing* that K is reliable.

But suppose that I know via vision that I'm seeing a red table, and that I know that this, at least in conjunction with the table's looking red, counts in favor of vision's reliability. It does not follow from these two suppositions, however, that I know that vision is reliable. In fact, this is precisely the point of Markie's argument for the claim that non-BKS theories are susceptible to the problem of easy knowledge.

Still, there are cases in which the following generalized version of (IP^*_K) holds:

(IAP^*_K) If I know via K that p , and if I know that p counts in favor of or entails q ,
then my belief that p supports my knowing that q .

The cases in which (IAP^*_K) holds are instructive. Consider a case in which I know via vision that the table is red, and in which I know that the table's being red entails that it isn't green. In ordinary cases, in which the proposition, *the table is green*, is neither a skeptical hypothesis nor a conjunctive component of such a hypothesis, it is entirely legitimate to say that I know that the table isn't green.

What's the difference, then, between this case and the case involving the reliability of vision, which shows that (IP^*_K) is false? The difference, I suggest, is just the difference between how I can come reasonably to believe, on one hand, that the table isn't green and how I can come reasonably to believe, on the other hand, that vision is reliable. In the former case, it seems perfectly clear that vision and reason alone will allow me reasonably to believe that the table isn't green. Yet vision and reason alone won't allow me reasonably to believe that vision itself is reliable. As I note in the next section, if I'm reasonably to hold a belief concerning vision's reliability—for example, that something's looking red accurately indicates its being red—I must rely not only on reason and vision, but also on my inductive evidence for the reliability of vision,

which will likely involve perception, memory and testimony. What shows (IAP*_K) to be false are cases in which S knows via K that *p* and via K⁺ that *p* entails *q*, but in which she can reasonably believe that *q* only via some mechanism(s) other than—or in addition to—K or K⁺. (Note here that K⁺ need not be distinct from K, and that the ‘or’ in ‘K or K⁺’, as well as those used below in similar contexts, is the *inclusive* ‘or’.) To construct a true principle, then, we might restrict any particular application of (IAP*_K) to a single mechanism or to a single set of mechanisms. Here’s a suggestion:

(SSC) If S knows via K that *p*, and if S knows via K⁺ that *p* entails *q*, and if K or K⁺ will allow S reasonably to believe that *q*, then S knows that *q*.

We have derived this principle—*Single Source Closure*—from plausible epistemic principles and intuitions. This gives us a strong reason to accept it.

Moreover, (SSC) explains why we don’t have knowledge in cases in which knowledge seems to come too easily: While (SSC) does *not* allow us to know too easily in such cases, every one of the most plausible available closure principles *does*. Consider three such principles:

- (STR) For all epistemic agents S and propositions *p* and *q*, if S knows that *p* and that *q* is incompatible with S’s knowing that *p*, then S knows that $\sim q$.¹⁷
- (KLE) If *x* is known by S and *x* entails *y*, then S is in a position to know that *y*.¹⁸
- (HAW) Necessarily, if S knows *p*, competently deduces *q* and thereby comes to believe *q*, while retaining knowledge of *p* throughout, then S knows *q*.¹⁹

Suppose that S knows that the table is red. Suppose too that she competently deduces that the table is not white and illuminated by red lights, where she counts as performing a competent

¹⁷ See B. Stroud, *The Significance of Philosophical Scepticism* (Oxford UP, 1984), pp. 24, 27.

¹⁸ See Klein, ‘Closure Matters: Academic Skepticism and Easy Knowledge’, p. 167.

¹⁹ See J. Hawthorne, *Knowledge and Lotteries* (Oxford UP, 2004), p. 82. Compare T. Williamson, *Knowledge and Its Limits* (Oxford UP, 2000), p. 117.

deduction at least in part because she knows that *the table is red* entails *the table is not white and illuminated by red lights*, and because she knows that the table's being white and illuminated by red lights is incompatible with her knowing that the table is red. Finally, suppose that S comes to believe that the table is not white and illuminated by red lights on the basis of her competent deduction, and that she knows throughout that the table is red. In this case, (STR), (KLE) and (HAW) are all forced to maintain that S knows—too easily, it seems—that the white-table hypothesis is false. And I submit that these are the most plausible closure principles on the market. Yet, as we will see in the next section, (SSC) need not maintain that S knows that the white-table hypothesis is false. This gives us reason to prefer (SSC), as well as to think that it affords an adequate explanation of why we don't have knowledge in cases in which knowledge seems to come too easily.

4. SOLVING THE PROBLEM OF EASY KNOWLEDGE WITH SINGLE SOURCE CLOSURE

We can now see that Single Source Closure provides no bridge to easy knowledge. Let's begin with the closure version of the problem, which we can illustrate with the following claims:

- (4) I know that the table is red.
- (5) I know that the table's being red entails that it isn't white and illuminated by red lights.
- (6) If I know that the table is red, and that the table's being red entails that it isn't white and illuminated by red lights, then I know that the table isn't white and illuminated by red lights.

Given (6), there's no reason why it wouldn't follow from (4) and (5) that

- (7) I know that the table isn't white and illuminated by red lights.

But the closure principle of which (6) is an instance is not the most plausible one available to us. When we opt instead for an instance of (SSC), we get

- (8) If I know via vision that the table is red, and if I know via reason that the table's being red entails that it isn't white and illuminated by red lights, and if vision or reason will allow me reasonably to believe that the table isn't white and illuminated by red lights, then I know that the table isn't white and illuminated by red lights.

Given (8), the argument for (7) must include the following premise if it's to be valid:

- (9) Vision or reason will allow me reasonably to believe that the table isn't white and illuminated by red lights.

(9), however, is false. We've already seen, in the second paragraph of Section 3, that since scenarios in which the white-table hypothesis is false are *ex hypothesi* visually indistinguishable from scenarios in which it is true, vision alone doesn't support our belief that the white-table hypothesis is false. And reason, on its own, also seems ill-equipped to support it—since the belief concerns the color of the table, it seems that evidence in its favor should come ultimately from vision. Moreover, since vision cannot distinguish between the table's being red and its being white and illuminated by red lights, reason, even working in conjunction with vision, gives us grounds neither for accepting nor for denying the white-table hypothesis. So, since (9) is an essential component of any valid argument from (4), (5) and (8) to (7), closure—that is, Single Source Closure—does not support my knowing (too easily) that the table isn't white and illuminated by red lights. By demanding an additional—and, as it turns out, false—premise in the argument for (7), closure, properly understood as Single Source Closure, takes away the bridge from (4) to (7).

What, then, will allow me reasonably to believe that the white-table hypothesis is false? My *inductive* evidence, as Cohen himself suggests,²⁰ is very much involved here. Of course, my inductive evidence—much of it, at least—is ultimately visual. For example, I’ve never seen anyone perpetrate any such deception; or noticed light fixtures, equipped with red bulbs, hanging above a white table in a furniture store; or been in someone’s home and on placing my hand between the ceiling and a white table, discovered that it took on a reddish hue. Still, some of my inductive evidence derives from sources other than vision. For example, I have never, in any place, overheard someone claim that he has been tricked into thinking that the white tables around here are red; I have never been in a furniture store where I noticed teenaged employees snickering as I describe a particular table as red; and so forth.

Moreover, my inductive evidence is constituted in part by *testimonial* evidence. For example, no report that I’ve ever received, through whatever medium, about visual experiences of tables at furniture stores, offices, or homes is a report of someone’s having been tricked, in the way described in the white-table hypothesis, into believing that a white table was red.

Memory also plays a role in justifying my belief that the white-table hypothesis is false, for I recall never having tricked anyone, never having been tricked myself, and never having heard of anyone being tricked in the way described in that hypothesis. It seems evident, then, that more than just vision or reason is involved in justifying my belief that the white-table hypothesis is false.

Turning to the bootstrapping and non-BKS versions of the problem of easy knowledge, we find in each of these cases too that (SSC) does not support our knowing too easily. First, the connections required for bootstrapping, and which when multiplied “by taking a few more looks”

²⁰ See BKPEK, p. 313, and ‘Contextualism, Skepticism, and the Structure of Reasons’, p. 74.

seem to support my knowing too easily that something's looking red is a reliable indication of its being red, take something like the following form:

- (10) The table is red.
- (11) The table looks red.
- (12) Therefore, the table's looking red is an accurate indication—that is, it indicates in this particular case—that the table is red.

Yet if knowledge of (10) and (11) supports knowledge of (12), there will be a closure principle that specifies how (10) and (11) are related to (12). Opting for (SSC), we get

- (13) If I know via vision that the table is red and via introspection that the table looks red, and if I know via reason that its both looking red and being red entails that the table's looking red is an accurate indication of its being red, and if vision, introspection or reason will allow me reasonably to believe that the table's looking red is an accurate indication of its being red, then I know that the table's looking red is an accurate indication of its being red.

Here, (SSC) necessitates the addition of the following premise:

- (14) Vision, introspection or reason will allow me reasonably to believe that the table's looking red is an accurate indication of its being red.

Yet (14) is false. Cases in which vision and introspection agree, as when I believe via vision that the table is red and via introspection that the table looks red, are the best candidates for cases in which vision or introspection allow me reasonably to believe that the table's looking red accurately indicates its being red. Notice, however, that the table's looking red is an essential part of the process that produces my belief that the table is red; when my belief is produced by

vision alone, if the table didn't *look* red to me, I wouldn't believe that it *is* red. The question here, then, is whether *vision* accurately indicates that the table is red. Intuitively, to have reason to believe that it does, I may not rely solely on what introspection, since in cases like this it reports simply on that part of the visual process that is constituted by the way things look to me, and vision itself have to say—I need to have received, in at least *some* cases in which I have employed vision, extra-visual confirmation of vision's reports. No amount of information that I receive from vision alone can give me reason to believe that vision's reports are true to the (extra-visual) facts.

Reason fails to come to the rescue here. For one thing, reason alone simply isn't equipped to confirm that the table is red. Moreover, reason alone, rather than justifying the belief that the table's looking red accurately indicates its being red, reminds us that the table can look red even when it *isn't* red, and that this can be true even in situations in which everything seems perfectly fine. (14), then, is false.

What *will* allow me reasonably to believe that the table's looking red is an accurate indication of its being red? I need to draw on my inductive evidence, on my memorial evidence and on my testimonial evidence. First, my inductive evidence, some of which is culled from remembered cases, supports the belief that the table's looking red accurately indicates its being red. Cases that are similar to this one in the relevant respects are cases in which

- (a) I take myself to know via vision that some object, *o*, is of a particular color, *c*,
- (b) I take myself to know via introspection that *o* looks *c*,
- (c) I take myself to know via reason that *o*'s both looking *c* and being *c* entails that *o*'s looking *c* accurately indicates its being *c*,
- (d) nothing suggests to me that there is anything wrong either with my vision, or with

my introspective capacity, or with my capacity to reason, and

- (e) nothing suggests to me either that *o* might not in fact be *c*, or that *o* might not in fact look *c*, or that *o*'s both looking *c* and being *c* might not in fact entail that *o*'s looking *c* accurately indicates its being *c*.

In the cases that I recall that have these features—and that are therefore like the present case in the relevant respects—I have had ample reason to believe that something's looking *c* accurately indicated its being *c*. For when I acted on that belief, I received every indication that it was true. When, for example, I requested the *c*-colored object, those around me supplied the object I requested, which looked *c* to me. Or when I reported that the object was *c*, those around me agreed with my reports, whether openly or tacitly. And when conceptually competent others reported that some object was *c*, I noticed that the object they indicated looked *c* to me. So, given that these remembered cases are similar in the relevant respects to the present case, I have some inductive support for the claim that the table's looking red accurately indicates its being red.

Testimony also seems important here in at least three ways. First, we have just seen that the testimony of others, in past cases that are similar to the present one, is a source of evidence for the belief that something's looking *c* accurately indicates its being *c*. Second, others report experiences that are consonant in this regard with my experiences: So far as I can recall, people always report—or they report in the extraordinarily vast majority of cases—that in cases that are similar in the relevant respects to the present one, they received every indication that something's looking *c* accurately indicated its being *c*. Finally, the testimony of others can support my belief in *this* case, just as it supports my belief in similar past cases. Each of these things, in addition to those discussed in the preceding paragraph, is an important source of

inductive evidence for my belief. Thus, more than just vision, introspection or reason is involved in justifying my belief that the table's looking red accurately indicates its being red.

Let's turn finally to the non-BKS version of the problem of easy knowledge, which seems to arise from connections such as the following:

- (15) '[My inductive evidence for the reliability of color vision]' (WBKEK, p. 421)
- (16) 'Color perception of the sort I'm employing right now is reliable and the table looks red' (WBKEK, p. 421).
- (17) I know that the table is red.
- (18) I know that the table's being red entails that it isn't white and illuminated by red lights.
- (19) If I know that the table is red and that its being red entails that it isn't white and illuminated by red lights, then I know that it isn't white and illuminated by red lights.
- (20) Thus, I know that it isn't white and illuminated by red lights.

To see that we need not countenance this connection, we can replace (19) with this instance of (SSC):

- (8) If I know via vision that the table is red, and if I know via reason that the table's being red entails that it isn't white and illuminated by red lights, and if vision or reason will allow me reasonably to believe that the table isn't white and illuminated by red lights, then I know that the table isn't white and illuminated by red lights.

Once again, (SSC) necessitates the addition of (9), which, as we have seen, is false. (SSC) therefore fails to license the above connection, which supports easy knowledge of the fact that the white-table hypothesis is false.

In conclusion, I point out that my solution to the problem of easy knowledge does several things that we want such a solution to do. First, Single Source Closure helps us to see why the connection between (10) and (11) seems *not* to support my knowing that something's seeming to have a certain feature accurately indicates its having that feature, but why the connection between (1) and (2) *does* seem to support such knowledge. (SSC) supports our intuitions in both of these cases, providing no bridge from (10) and (11) to (12), but maintaining that there is a path from (1) and (2) to (3), a path that runs through

(21) Vision, reason or my somatosensory faulty will allow me reasonably to believe that its feeling as if I've just cut my finger accurately indicates that I've just cut my finger.

(21), unlike (14), is entirely reasonable. (14) is false at least in part because one of the processes it mentions—namely, introspection—is used to determine whether I possess certain qualitative states, where those states are a part of another process mentioned in (14), namely, vision. To put this another way, one of (14)'s processes is used to monitor or to examine (a part of) another, and so those two processes fail to operate sufficiently independently. But introspection, when it's used in this way, cannot contribute to our knowing that the table's looking red accurately indicates its being red, for it accounts only for a part of the visual process, namely, its qualitative upshots. And in general, the fact that a *part* of a process seems upon examination to proceed successfully does not accurately indicate that the process *as a whole* proceeds successfully.

But none of the processes mentioned in (21) is used to monitor another, that is, to determine whether I possess certain states that are a part of another process. Vision and somatosensation operate independently in (21), which means, among other things, that neither vision nor somatosensation is used to determine whether I possess the qualitative states associated with the other process. In this case, my somatosensory faculty makes me feel as if the world has certain features, and vision reports on *just those features*. Vision files its report from *outside* my somatosensory faculty, as it were. This suggests that vision can contribute to my knowing that its feeling as if I've just cut my finger accurately indicates that I've just cut my finger, for those features of the world to which it responds, independently of my somatosensory faculty, can confirm the feeling that is generated by somatosensation. This helps to explain why (1) and (2) together seem to support (3).

Now, nothing in Section 3's formulation of (SSC) suggests that (1) and (2) will support (3), but that (10) and (11) will fail to support (12). Nevertheless, when we compared these two cases, we found that (SSC) necessitated the introduction of (14) and (21), respectively, and this put us in position to discover an epistemologically significant difference between the two cases. Our comparison of the two cases, when done in the light of (SSC), suggests that it is a part of our epistemological makeup to hold that a source, K_1 , cannot support our knowing that the qualitative states associated with another source, K_2 , accurately indicate the truth if K_1 acts only to account for those qualitative states. We can now reformulate (SSC) so that it accounts for this discovery:

(SSC*) If S knows via K that p_1, p_2, \dots, p_n , and if S knows via K^+ that p_1, p_2, \dots, p_n entails q , and if K or K^+ will allow S reasonably to believe that q , then S knows that q , where if K is a set of two belief-producing mechanisms,

neither member of the set acts only to account for the qualitative states associated with the other member of the set.

While (SSC*) certainly applies when q concerns whether something's seeming to have a certain feature accurately or reliably indicates its having that feature, it also applies more generally.

Suppose that I know via vision that Lester is in the apartment, and that I know via audition that Myrtle is in the apartment. Suppose further that I know via reason that if Lester is in the apartment and Myrtle is in the apartment, then they are both in the apartment. Since vision, audition or reason—in fact, given the right circumstances, since vision or audition alone—will allow me reasonably to believe that both Lester and Myrtle are in the apartment, (SSC*) says that I know that they are both in the apartment. And this is so at least in part because vision does not act only to account for audition's qualitative states, and audition does not act only to account for vision's qualitative states.

Second, my solution makes a place for my inductive evidence for the falsity of certain skeptical hypotheses. According to Cohen, the problem of easy knowledge is generated at least in part because 'my inductive evidence against the possibility that there are red lights shining on the table [for example] turns out to be irrelevant to my knowing the table is not white with red lights shining on it' (BKPEK, p. 313). But on a solution rooted in (SSC), my inductive evidence for the belief that the table isn't white and illuminated by red lights is not irrelevant—and, in fact, it can significantly contribute—to my knowing that the table isn't white and illuminated by red lights.

Third, my solution 'explain[s] why it is that, for any particular piece of basic knowledge [for example, my knowledge that the table is red], it seems that we can inferentially expand it in some ways but not in others' (Neta, 'A Contextualist Solution to the Problem of Easy

Knowledge', p. 196). For example, it seems perfectly acceptable to maintain that I can know, on the basis of my knowledge that the table is red, that it's not green, or that it will match the carpets in my living room. Single Source Closure explains why we find this acceptable: When we know via vision that the table is red and via reason that its being red entails that it isn't green, then according to (SSC), since vision or reason can justify our belief that the table isn't green, we know that it isn't green. Single Source Closure is therefore entirely consistent with our being able to expand our knowledge inferentially.

Yet, as we have seen, it seems that I don't know, on the basis of my knowledge that the table is red, that it isn't white and illuminated by red lights. Once again, Single Source Closure helps to explain why this is so. From the truth of Single Source Closure and the fact that I know via vision that the table is red, it does not follow that I know that the table isn't white and illuminated by red lights. Moreover, we've seen that (9), which is the relevant instance of one of the conditions specified in the antecedent of Single Source Closure, is false. Given this, Single Source Closure provides no bridge from my knowledge that the table is red to my knowledge that it isn't white and illuminated by red lights. In this way, (SSC) helps to explain why we can inferentially expand our knowledge in some ways but not in others.

Fourth, my solution requires neither that BKS theories distinguish between animal knowledge and reflective knowledge, nor that non-BKS theories adopt contextualism. Understanding closure as Single Source Closure, we can solve the problem of easy knowledge without altering or supplementing any epistemological theory, and we learn something important about how certain pieces of knowledge are related: The connection between knowing that p and knowing that q includes the condition that my belief that q be justifiable by the source that justifies my belief that p or by the source that justifies my belief that p entails q .

Finally, closure principles are meant in part, it seems, to capture the intuition that when S's knowledge that p and her knowledge that p entails q are sufficient for her knowing that q , then her evidence for p and for p entails q is *all it takes* for her to know that q . Single Source Closure is perhaps the only closure principle on the market that respects the limits of this intuition. Other closure principles, including those discussed at the end of Section 3, run roughshod over this intuition, allowing evidence for p and for p entails q to be sufficient for knowledge that q even when such evidence is clearly *not* sufficient for knowing that q . And this, of course, gives rise to the problem of easy knowledge. The suggestion here, then, is that the source of the problem of easy knowledge is not to be found in any epistemological theory, but in a certain way of formulating epistemic closure principles. Although this is not the place to pursue these matters, it strikes me that this might also be the source of other epistemological confusions, and so the lessons learned here might help us to solve not only the problem of easy knowledge, but other worrisome epistemological problems as well.²¹

²¹ For their very helpful comments, I thank Albert Casullo, Leo Iacono, Peter Kung, Peter Murphy, Takashi Yagisawa, Masahiro Yamada and an anonymous referee.