

A CONTEXTUALIST SOLUTION TO THE PROBLEM
OF EASY KNOWLEDGE

Ram NETA
University of North Carolina

Summary

Many philosophers hold some version of the doctrine of “basic knowledge”. According to this doctrine, it’s possible for S to know that p, even if S doesn’t know the source of her knowledge that p to be reliable or trustworthy. Stewart Cohen has recently argued that this doctrine confronts the problem of easy knowledge. I defend basic knowledge against this criticism, by providing a contextualist solution to the problem of easy knowledge.

I. *Basic knowledge and the problem of easy knowledge*

Many philosophers claim that situations of the following kind are possible:

S knows that p, and S obtains this knowledge by means of source X, but S doesn’t know that X is reliable or trustworthy.¹

For example, a child might know that his mother is home, and he might know this by seeing his mother at home, but all of this can be true even if he doesn’t know that his vision is reliable or trustworthy. Again, I might know by means of testimony and memory that I was born on a certain date, even if I don’t know that the testimony or the memory on which this knowledge is based is reliable or trustworthy. More exotically, an amnesiac might see a tree in front of her, and so come to know that there is a tree in front of her, even if she doesn’t recall any experiences on the basis of which she could come to know that her visual perception is at all reliable or trustworthy.

1. Virtually all externalists believe this, and so do many internalists. The view is criticized by Sellars (1963) and BonJour (1978), but those criticisms are rebutted in Alston (1983).

I will follow Cohen (2002) in using the phrase “basic knowledge” to denote all such knowledge—i.e., all knowledge possessed without knowledge of the reliability or trustworthiness of its source. What I’ll call the “doctrine of basic knowledge” is simply the view that we can possess basic knowledge, i.e. we can possess knowledge without also knowing its source to be reliable or trustworthy. In other words:

(The Doctrine of Basic Knowledge) For some subject S, some source x, and some proposition p: S knows that p, by means of source x, and S doesn’t know that x is reliable or trustworthy.

The doctrine of basic knowledge is widely (but not universally) accepted among epistemologists.

Recently, different versions of the doctrine of basic knowledge have been subject to a particular objection—roughly, that they do not have the resources to explain why it’s impossible to obtain knowledge by means of certain kinds of inferences, which I will henceforth call “easy knowledge” inferences. (I’ll characterize these inferences below.) Cohen (2002) refers to this objection, in its most general form, as “the problem of easy knowledge”.² In this paper, I defend the doctrine of basic knowledge (the various versions of which I will henceforth refer to as “basic knowledge theories”) from this objection. More specifically, I offer a contextualist solution to the problem of easy knowledge.

Many philosophers have already proposed solutions to the problem of easy knowledge.³ Almost all of these proposed solutions reject a presupposition of the problem, namely, that we generally cannot gain knowledge that p by means of an easy knowledge inference to the conclusion that p. But I think that we should not simply reject this presupposition. As I argue elsewhere, if we simply reject this presupposition, then we will not be able to distinguish, in a principled way, easy knowledge inferences from certain other inferences that are not (on anybody’s view) warrant-transmitting.⁴ My contextualist solution to the problem of easy knowledge does not give up the doctrine of basic knowledge,

2. Fumerton (1995, 175-6) and Vogel (2000) both level this line of criticism against reliabilism.

3. See, for instance, Van Cleve (2003), Davies (2003), Klein (2004), Pryor (2004), Markie (2005), and Bergmann (2004).

4. I argue for this in Neta (*forthcoming b*).

and does not simply reject the presupposition that the easy knowledge inferences cannot transmit warrant from premises to conclusion.

Before setting out the problem of easy knowledge, we must first get clearer about what the doctrine of basic knowledge says. In stating the doctrine of basic knowledge, I followed common practice in talking about “sources” of knowledge. But what is a “source” of knowledge? Some philosophers use the term “source” to refer to the cognitive faculties that produce knowledge, or to the evidence on which knowledge is based. But what, if anything, is the genus of which cognitive faculties and evidence are two species? Different proponents of basic knowledge theories might offer different answers to this question. For the purposes of stating the doctrine of basic knowledge in order to articulate the problem of easy knowledge, we can lay down the following general constraint on what we will count as a “source” of knowledge:

If X is a “source” of S’s knowledge that p, then all of the following conditions obtain:

- (i) S believes that p, and
- (ii) X is a reason that S has for believing that p, and
- (iii) S believes that p for, at least, that very reason (X)⁵, and
- (iv) X is a token of a type T1, and p a token of type T2, such that, in general, tokens of T1 are reasons for people to believe propositions of type T2, and
- (v) S can know that (i)–(iv) are all true, without X being a source of that knowledge.

I’ll summarize conditions (i)–(iii) by saying that X is “a reason why” S believes that p. I can then say that X is a source of S’s knowledge that p only if X is a reason why S believes that p, and things like X are reasons for people to believe things like p, and finally S can know all of this without X being a source of *that* knowledge. When I discuss the problem of easy knowledge below, it will become clear why it is that, for the purposes of explaining the doctrine of basic knowledge in a way that will make clear why it seems to be subject to the problem of easy knowledge, we should understand a “source” of knowledge as something that satisfies the five conditions listed above.

5. S’s believing that p for that very reason is compatible with S’s believing that p for other reasons as well—our believings are sometimes overdetermined.

To illustrate: If Fred's testimony is a source of Jones's knowledge that Mabel will be late to the party, then

(P1) Fred's testimony is a reason why Jones believes that Mabel will be late to the party.

And Jones can know that P1 is true, without Fred's testimony being Jones's reason to believe that P1 is true. (Presumably, if Jones believes that P1 is true, that is because Jones can recognize, upon reflection, what reasons move him to believe that Mabel will be late to the party.)

Again: If the smell of the stew is a source of Smith's knowledge that there's too much paprika in the stew, then

(P2) The smell of the stew is a reason why Smith believes that there's too much paprika in the stew.

And Smith can know that P2 is true, without the smell of the stew being Smith's reason to believe that P2 is true. (Once again, if Smith believes that P2 is true, that is because Smith can recognize, upon reflection, what reasons move him to believe that the stew contains too much paprika.)

But contrast: If reasoning is a source of my knowledge that $23 + 57 = 80$, then

(P3) My reasoning is a reason why I believe that $23 + 57 = 80$.

And I can know that P3 is true, without reasoning being my reason to believe that P3 is true. Since, in the standard case, I will not know that P3 is true without reasoning being a reason why I believe that P3 is true, it follows that, in the standard case, reasoning does not satisfy the preceding characterization of a "source" of knowledge. For the purposes of this problem, we must understand sources to be individuated more finely.

When we state the problem of easy knowledge below, it will become clear that, if an epistemologist does not think of sources as satisfying the complicated constraint that I've just laid down, then her view will not be even apparently subject to the problem of easy knowledge. We should therefore understand the problem of easy knowledge as a problem that is posed for basic knowledge theories that conceive of sources as satisfying the complicated constraint above.

Now that we've clarified the doctrine of basic knowledge, what is the problem of easy knowledge? Cohen presents it as having two versions, which I consider in sections II and III below.

II. *The first problem of easy knowledge: one-premise deductions against defeaters*

The first problem of easy knowledge is generated by the following plausible epistemic closure principle:

One-Premise Closure: If S knows that p, and S competently deduces q from p without losing her knowledge that p, then S knows that q.⁶

Notice the qualification that S retain her knowledge that p while deducing that q. This is important. If, while performing the deductive inference, S comes to lose confidence in p, or comes to lose her justification for believing that p, she will not be able to earn knowledge that q on the basis of the deduction. This point is distinct from the point that S's deduction must be "competent". S could retain her knowledge that p while making the deduction, but if S bungles the deduction in some way then, even if she somehow still ends up at the conclusion that q, the conclusion doesn't count as knowledge. Bungled deductions do not transmit knowledge from premise to conclusion—not even if the conclusion reached happens to be one that actually follows from the premise.⁷

Now, how does One-Premise Closure pose a problem for basic knowledge theories? Let's consider a case that Cohen describes:

'Suppose my son wants to buy a red table for his room. We go in the store and I say, "That table is red. I'll buy it for you." ... he worries, "Daddy,

6. Here I state the version of Closure similar to one recently defended by Hawthorne (2004). (Also see Williamson (2000, 117)) It differs from the version of Closure put forward in Cohen (2002), but I think that the problems of easy knowledge can be more clearly stated using Hawthorne's version of Closure than Cohen's own version.

7. Harman (1973, 157) questions whether there is any such psychological act as deductive inference. He points out that, when we construct a deductive argument, we typically do not perform any psychological act of deductive inference. He is right about that. But it's also true that, when we follow a deductive argument that's presented to us, we typically do perform a psychological act of deductive inference.

what if it's white with red lights shining on it?" I reply, "Don't worry—you see ... it is red, so it's not white but illuminated by red lights." Surely he should not be satisfied with this response. Moreover I don't think it would help to add, "Now I'm not claiming that there are no red lights shining on the table, all I'm claiming is that the table is not white with red lights shining on it."⁸

Now, is Cohen right to claim that the reasoning of the father in his example (whom I'll henceforth call "Stewart") is bad reasoning—more specifically, that such reasoning does not provide Stewart with knowledge that the table is not white with red lights shining on it? To answer this question, let's consider how Stewart could have achieved his knowledge that the table is red. Suppose that Stewart achieved knowledge that the table is red by means of the following deductive inference:

The table is not white with red lights shining on it.
If the table is not white with red lights shining on it, then it is red.

The table is red.

If Stewart achieved his knowledge that the table is red by means of this deductive inference, and he achieved his knowledge of the premises of the inference in some other way altogether, then the problem with the reasoning described in Cohen's example is that such reasoning cannot provide Stewart with knowledge that the table is not white with red lights shining on it, since Stewart had knowledge of that fact already. (And if he lost that bit of knowledge while making the deduction that Cohen describes, then the deduction won't give him knowledge of the conclusion, since it rests on a premise that he doesn't know to be true.) More generally, if Stewart achieved his knowledge that the table is red on the basis of his knowledge that the table is not white with red lights shining on it, then of course the reasoning described in Cohen's example cannot be the way that Stewart acquires knowledge that the table is not white with red lights shining on it.

If, however, Stewart achieved his knowledge that the table is red in some other way—for instance, simply by virtue of seeing the table looking red in normal circumstances—then again, it seems that Stewart cannot acquire, merely by means of the reasoning described, knowledge

8. Cohen (2002).

that the table is not white with red lights shining on it. For, if all he has to go on is that the table looks red, then he cannot achieve knowledge that the table is not white with red lights shining on it.

But if Stewart has basic knowledge that the table is red, and One-Premise Closure is true, and he can competently deduce the proposition “the table is not white with red lights shining on it” from the proposition “the table is red” without losing his knowledge of the latter, then this reasoning can provide him with knowledge that the table is not white with red lights shining on it. So something is wrong here: it seems that (a) Stewart’s reasoning in Cohen’s example is actually good, knowledge-transmitting reasoning, or (b) One-Premise Closure is not true, or (c) Stewart cannot competently deduce the proposition “the table is not white with red lights shining on it” from the proposition “the table is red” without losing his knowledge of the latter, or (d) Stewart cannot have basic knowledge that the table is red. Since the first three of these three options are extremely implausible, we are under considerable pressure to adopt the fourth option, and so deny that Stewart can have basic knowledge that the table is red. But this example is perfectly representative: if Stewart in Cohen’s example cannot have basic knowledge that the table is red, then none of us can ever have basic knowledge of any particular contingent fact about the external world.

So the first problem of easy knowledge that besets basic knowledge theories is illustrated by this example: If Stewart has basic knowledge that the table is red, then it seems that he should be able to know by deduction, all too easily, that the table is not white with red lights shining on it.

Now, in this example, a source of Stewart’s basic knowledge that the table is red is this: the table looks red. So our example suggests a general problem of the following form: if S has basic knowledge that p, and One-Premise Closure is true, then it follows that S can, by means of a simple and obvious one-step deduction, come to know—all too easily, it seems—that particular defeaters do not obtain (e.g., that it merely appears to her as if p). Why is it that, in the relevant range of cases, S cannot gain such knowledge so easily? The thought seems to be this: in the relevant cases, the reason why S believes that p—i.e., the source of S’s knowledge that p—cannot serve as a sufficient reason for S to believe that certain defeaters don’t obtain. So, if we accept One-Premise Closure, it seems that we’re driven to the following conclusion: if S knows that p, and r is a reason why S believes that p, then S must also

know that *r* is a reliable or trustworthy reason. If *S* doesn't also know that *r* is reliable or trustworthy, then *S* will be able to gain knowledge by deduction from *p*—and all too easily—that certain defeaters of *r* don't obtain. If *r* does not supply *S* with such knowledge all by itself, then it seems *S* cannot acquire such knowledge merely by deduction from things (like *p*) that she knows solely on the basis of *r*. And if *r* does supply *S* with such knowledge, then *S*'s knowledge that *p* is not basic. In other words, if we accept One-Premise Closure, then we must reject any doctrine of basic knowledge that conceives of the sources of our knowledge as satisfying constraints (i), (ii), and (iii) listed above in our account of what's involved in something's being a "source" of knowledge. In the next section, we'll see why constraints (iv) and (v) are important also.

III. *The Second Problem of Easy Knowledge: Bootstrapping Inferences Against Counterpossibilities*

The second problem of easy knowledge is generated by assuming the following epistemic closure principle:

Ampliative Closure: At least some ampliative arguments are such that, if *S* knows all of the premises to be true, and she competently infers the conclusion from the premises without losing her knowledge that the premises are true, then she knows the conclusion to be true.

Ampliative Closure makes an existence claim, but it doesn't specify which ampliative arguments preserve knowledge in this way—that is a substantive and contentious question within confirmation theory. Ampliative Closure says only that there are some such ampliative arguments.

Now, how does this pose a problem for basic knowledge doctrines? Consider an ampliative argument which exemplifies the property mentioned in Ampliative Closure—someone can achieve knowledge of its conclusion by competently inferring it from the premises. Let's suppose that the following inductive argument has this property:

- (1a) At t1, it looks to me as if the table is red.
- (1b) At t1, the table is red.
- (2a) At t2, it looks to me as if the cat is on the mat.
- (2b) At t2, the cat is on the mat.
- (3a) At t3, it looks to me as if it's snowing.
- (3b) At t3, it's snowing.
- ...

(conclusion) The way things look to me is (generally) the way they are.

In this inference, let's suppose, each (a) premise states a source of the knowledge expressed by the (b) premise that immediately follows it. And let's suppose that someone knows all of these premises to be true. More specifically, let's suppose that someone—call her “Nancy”—has basic knowledge of each of the (b) premises. (If, for whatever reason, you are inclined to think that this cannot be reasonably supposed of the (b) premises I've chosen, then imagine other (b) premises of which it can be reasonably supposed. If you accept the doctrine of basic knowledge, then you are committed to believing that there are *some* such premises.) Since each (a) premise states a source of the basic knowledge expressed by the corresponding (b) premise, Nancy is able to know the (a) premises to be true without reliance on those sources themselves. (In the case I've chosen, Nancy can—I assume—introspectively know the (a) premises to be true.) Thus, Nancy knows (the truth of) all of the premises of this inductive argument. And the inductive argument for the conclusion is a straightforward case of induction by enumeration. Of course, some inductions do not give Nancy knowledge of their conclusion because those inductions rely on a small or unrepresentative data set. But this need not be a problem for the inductive argument above, since we can expand our data set to include everything that we know by means of the source in question (here, the source would be the looks of things).

What Nancy has here, then, is a seemingly good inductive argument from known premises to a conclusion. If she competently goes through this reasoning without losing her knowledge of the premises, can doing so give her knowledge of its conclusion? It seems not. If Nancy is wondering why she should trust her visual experiences, and she then goes through the reasoning above, she should not regard this reasoning as

supplying her with knowledge that she should trust her visual experience. But how could that be? It could be that induction by enumeration generally doesn't give us knowledge of its conclusion. But if this is so, then what sort of ampliative inference does give us knowledge of its conclusion? Whatever form that is, couldn't Nancy duplicate the problematic reasoning by means of that form of inference?

Whatever the answer to this last question, it needn't detain us. For even if the basic knowledge theorist isn't willing to grant that induction can ever give us *knowledge* of its conclusion, she will still confront a version of the Bootstrapping problem so long as she accepts the following closure principle.

Ampliative Closure concerning epistemic reasons: At least some ampliative arguments are such that, if S knows all of the premises to be true, and she competently infers the conclusion from the premises without losing her knowledge that the premises are true, then she has good epistemic reason to believe the conclusion to be true.

Even if we don't accept the original version of Ampliative Closure, we should at least accept Ampliative Closure concerning epistemic reasons. But how does this principle apply to our example above? Can the inductive argument stated above ever give Nancy good reason to believe its conclusion? Again, it seems clear that it cannot do so. If Nancy is wondering what good reason she has to trust her visual experiences, and she goes through this reasoning, she should not regard this reasoning as supplying her with extra reason to trust, or to believe that she should trust, her visual experience.⁹ Her reasons for believing that her visual experience is trustworthy are not improved by virtue of her going through this reasoning.

So it seems that we have to deny that Nancy knows all of the premises of the induction to be true. But then which of the premises does she not know to be true? If she doesn't know the (a) premises, then she doesn't know how things look to her, and that is absurd. Of course she knows how things *look* to her! Even if she doesn't know everything there is to

9. Some philosophers would deny this claim. They would say that such inductive arguments can—at least in some cases—give us good reason to believe their conclusions. See, for instance, Alston (1986; 1993), Sosa (1997), Van Cleve (2003), and Bergmann (2004). I argue against this view in my (*forthcoming b*).

know about how things look to her, and even if she is sometimes wrong about how things look to her, still she does (or at least we can coherently and plausibly stipulate that she does) know a great deal about how things look to her. At least, she knows enough about how things look to her that she can use such knowledge in an inductive argument of the form just sketched. Is it that Nancy doesn't know the (b) premises then? If Nancy has any basic knowledge of visible things around her, then there must be some premises which are such that she can have basic knowledge, supplied by vision, that they are true. Whatever those premises are, let them be the (b) premises of our bootstrapping argument. In that case, Nancy will know the (b) premises of the argument to be true. If Nancy has any basic knowledge of visible things around her, then either she doesn't know the (a) premises of the inductive argument, or the inductive argument somehow gives her an additional reason to believe its conclusion. Both of these last two options are absurd, and so we are under pressure to deny that Nancy can have basic knowledge.¹⁰

The problem of easy knowledge that is illustrated by this example can be spelled out more generally as follows. Suppose that S has basic knowledge that p1, p2, p3, etc. Also, the source of S's basic knowledge that p1 is r1, the source of S's basic knowledge that p2 is r2, the source of S's basic knowledge that p3 is r3, and so on, where the various p's constitute a projectible type of proposition and the various r's constitute a projectible type of reason to believe a proposition. Now, since r1 is the source of S's basic knowledge that p1, then, given our account of what's involved in something's being a "source" of knowledge, S can know, without relying on r1, that r1 is the reason why she believes that p1. Furthermore, S can know, without relying on r2, that r2 is the reason why she believes that p2, and again she can know, without relying on r3, that r3 is the reason why she believes that p3, and so on. In that case, S will know all of the premises of an inductive argument of the following form:

10. The arguments provided by Fumerton (1995) and Vogel (2000) are similar to the argument of this section.

r1 is my reason for believing p1.
p1.
r2 is my reason for believing p2.
p2.
r3 is my reason for believing p3.
p3.
...

r's are generally reliable reasons for believing p's.

S will know all of the even-numbered premises by means of the respective r's. And S will know all of the odd-numbered premises in some other way than by means of those same respective r's. Thus, merely by having basic knowledge of the p's, and also knowledge of which r's supply her with reason to believe which p's, S will have good reason to believe, by some ampliative inference, that r's are generally reliable or trustworthy.

But it seems clear that S cannot have good reason to believe that r's are reliable or trustworthy merely by means of any such ampliative inference. If S knows that r1 is a reason why she believes p1, and she knows that r2 is a reason why she believes p2, and so on, then how can it be rational for her to gain confidence in the reliability or the trustworthiness of the r's by means of this inference from those very beliefs, in conjunction with her belief of the p's? If a "source" of knowledge satisfies conditions (i)–(v) listed above in our account of "sources", then—it seems clear—such ampliative inferences cannot suffice for S to have good reason to believe their conclusion.

IV. *Cohen's proposed solution to the problems of easy knowledge*

After posing these problems for Basic Knowledge theories, Cohen proposes the following solution: allow that there is some basic knowledge (what Cohen, following Ernest Sosa, calls "animal knowledge"), and that this basic knowledge can be inferentially expanded in *some* ways, but it cannot be inferentially expanded by the kinds of inference that are used in the One-Premise Closure and Bootstrapping reasoning above. Cohen thus denies that certain plausible general principles of knowledge-transmission hold for basic knowledge.

Cohen candidly admits that this solution is *ad hoc*. It is *ad hoc* because it is unable to explain, by appeal to any already agreed-upon principles, why it is that, for any particular piece of basic knowledge, we can inferentially expand it in some ways but not in others. To illustrate: if Stewart has basic knowledge that the table is red, then he can know by deduction that the table is not green. So why is it that he can know by deduction that the table is not green, but he cannot know by deduction that the table is not white with red lights shining on it? Why is it that Stewart can inferentially expand his knowledge that the table is red into knowledge that the table is not green, but he cannot inferentially expand his knowledge that the table is red into knowledge that the table is not white with red lights shining on it? Again, if Nancy has basic knowledge of all of the (b) premises of the induction above (“the cat is on the mat”, “it’s snowing”, “the table is red”), then she should be able to know by induction (given normal background knowledge) that if the cat jumps into the snow and then onto the table, the part of the table’s surface that’s underneath the cat’s paws will become wet. But how could Nancy acquire knowledge of this fact by induction, if she couldn’t acquire knowledge of the reliability of her visual faculties by induction? Once again, what’s the difference between the two cases?

Cohen suggests that, while basic knowledge can be inferentially expanded in some ways, it cannot be inferentially expanded to reach conclusions that are not, as Nozick (1981) would say, “sensitive”. That is, basic knowledge cannot be inferentially expanded to produce knowledge that *p*, for any *p* which is such that we would still believe it (and by the same method) if it were not true. This proposal might help to explain why Stewart can inferentially expand his basic knowledge that the table is red into knowledge that the table is not green, but cannot inferentially expand his knowledge that the table is red into knowledge that the table is not white with red lights shining on it. But this view leaves us with two unanswered questions.

First, suppose that Stewart competently performs the following inference:

The table is red.

Therefore, it’s not the case that I falsely believe that the table is red.

Here again, the conclusion is insensitive. But can't Stewart gain knowledge of this conclusion by means of competently performing this inference? The answer to this question is not a clear "no". It seems not as clearly wrong to say that Stewart gains knowledge of the conclusion above by competently performing the inference above, as it does to say that Stewart gains knowledge of the conclusion "the table is not white with red lights shining on it" by competently performing the inference from "the table is red". Since Cohen's view can't explain this difference, it leaves us with an unanswered question.

And there is a second unanswered question: even if sensitivity is relevant in the way that Cohen thinks it is, still, why is *that* difference—the difference between sensitive and insensitive conclusions—one that should make a difference to the possibility of inferentially expanding one's basic knowledge? Cohen's solution, again, doesn't answer this question.

So the gap in Cohen's proposed solution should teach us this: we want a solution to the problem of easy knowledge to explain why it is that, for any particular piece of basic knowledge, it seems that we can inferentially expand it in some ways but not in others—even when the form of inference is the same across these cases. Such a solution should not only accurately predict which cases fall on which side of that distinction, but it should also explain to us why those cases fall where they do.

V. *Rule of context-shifting*

We can develop a satisfying solution to the problem of easy knowledge if we approach these problems from a view of knowledge and justification that I've defended elsewhere.¹¹ For present purposes, I will simply state—without arguing for—a rule of context-shifting that forms part of my account.

I begin with some terminology. Consider any two hypotheses—h1 and h2—that can be about anything whatsoever. Let's say that h1 and h2 are "introspectively indistinguishable for S" just in case:

(a) If h1 were true, then there would be a certain probability distribution over the mental states that S could be in (i.e., there is a certain

11. See Neta (2002; 2003; 2004; *forthcoming c*).

probability that S would be in mental state M1, there is a certain probability that S would be in mental state M2, etc.), and

(b) If h2 were true, then there could be a different probability distribution over the mental states that S could be in (i.e., there is either a different probability that S would be in mental state M1, or a different probability that S would be in mental state M2, etc.), and

(c) For any two mental states m1 and m2 which are such that the probability of S being in one of them rather than the other would differ depending upon whether h1 or h2 were true, the difference between m1 and m2 is not introspectively available to S.¹²

To say that the difference between two mental states is not “introspectively available” to S is just to say that there is no difference in what it’s introspectively like for S to be in one rather than the other mental state. As far as S’s introspection reveals, the two mental states are indistinguishable.

In order to clarify this definition of “introspective indistinguishability for S”, I’ll offer some examples of pairs of hypotheses that are introspectively indistinguishable for me:

- (A1) I am now drinking Coke.
- (A2) It am now drinking something that looks and tastes exactly like Coke.
- (B1) The creature before me is experiencing pain.
- (B2) The creature before me is behaving exactly as if she’s experiencing pain.
- (C1) I have hands.
- (C2) I am a brain in a vat being neurochemically stimulated to have the experiences that I’m now having, and that lead me to believe that I have hands.
- (D1) There are 3 teaspoons of sugar in the coffee that I’m now drinking.
- (D2) There are 3.1 teaspoons of sugar in the coffee that I’m now drinking.

12. If we assume determinism, then we can formulate this definition of introspective indistinguishability much more simply, as follows: h1 and h2 are introspectively indistinguishable for S just in case

If h1 were the case, then S would be in m1,

If h2 were the case, then S would be in m2, and

The difference between m1 and m2 is not introspectively available to S.

Each of the four pairs above contains two hypotheses that are introspectively indistinguishable for me. The D-pair cases might be qualitatively different, but this qualitative difference is—I assume—so slight that it falls below the threshold of introspective discriminability. Of course, which discrimination tasks we are capable of performing depends in part upon the circumstances in which those tasks are presented to us. (For instance, while it might be impossible to discriminate *m1* from *m2* when they're considered by themselves, it might be possible to discriminate *m1* from *m2* when they're considered along with *m3*, which is more similar to *m2* than to *m1*.) But when I say that a difference between *m1* and *m2* is “introspectively unavailable” to S, I mean to imply that there are *no* circumstances in which S can perform the task of introspectively discriminating *m1* from *m2*.

Now, let's say that a hypothesis H is “an uneliminated counterpossibility” with respect to S's knowing that p at t just in case (i) H implies that S doesn't know that p at t and (ii) H and “S knows that p at t” are introspectively indistinguishable for S. (In other words, what it's introspectively like for S when H is true is the same as what it's introspectively like for S when “S knows that p at t” is true, and yet H is incompatible with “S knows that p at t”.) An appraiser X “raises” an uneliminated counterpossibility with respect to S's knowing that p at t just in case X (seriously and sincerely) treats that counterpossibility as a reason to regard S's epistemic state as being worse than it would be were it not for that uneliminated counterpossibility.

Now using the terminology just introduced we can state the following rule:

(R) When X raises a hypothesis H that is an uneliminated counterpossibility with respect to S's knowing that p at t, X moves into a context of ascription¹³ in which she can truthfully affirm, and cannot truthfully deny, that S's total evidence at t includes all and only the evidence that S would have at t whether or not H is true.¹⁴

13. It is somewhat misleading to say that X “moves” into such a context of appraisal, since X may already have been in such a context. Maybe it is more accurate to say that, by raising the counterpossibility, X sees to it that she is in such a context.

One's listeners may refuse to enter into one's context of appraisal by refusing to treat H as relevant to the appraisal of S's epistemic state. One can expect this to happen when one raises skeptical hypotheses in a courtroom or a scientific laboratory, for instance.

14. Some epistemological internalists may object that if H and “S knows that p at t” are

Since one's beliefs are justified by all and only one's evidence for them, it follows from (R) that ascriptions of justification are also semantically context-sensitive.

I have elsewhere argued¹⁵ that (R), in conjunction with a particular analysis of knowledge, provides us with an elegant, unified, and intuitively plausible solution to the Gettier problem, the lottery paradox, the dogmatism paradox, skeptical problems of underdetermination, the skeptical closure puzzle, and the problem of distinguishing epistemic from non-epistemic justification. Furthermore, an independently well-motivated account of the function of evidence ascriptions¹⁶ predicts that such ascriptions should be semantically context-sensitive in accordance with some rule like (R). It seems to me that the convergence of all of these apparently independent considerations provides a compelling case for (R). But I do not have the space to rehearse all of these independent considerations here. I mention them now solely in order to indicate the variety of grounds on which I've elsewhere defended (R). In this paper, I will add to those grounds by appealing to (R) in order to solve the problem of easy knowledge.

When we raise uneliminated counterpossibilities to S's knowing that p at t, we thereby (in accordance with rule (R)) move into a context of appraisal in which we cannot truthfully say, and can truthfully deny, that S has evidence for p (rather than q) at t. In that new context, S's belief that p cannot be truthfully said to be well justified at all, and so cannot truthfully be said to be knowledge. Thus, by raising skeptical counterpossibilities to S's knowing that p, we make it true to say that S's belief that p does not count as knowledge.

VI. *Solving closure and bootstrapping*

Let's now show how we can use the accounts of knowledge and justification developed above to solve the problems of easy knowledge.

introspectively indistinguishable for S, then what evidence S has cannot depend upon whether or not H is true—and this is true independently of any context of epistemic appraisal. This objection rests upon the assumption that the only facts that qualify as someone's evidence for p at t are facts that the agent can know by introspection at t. But, as Alvin Goldman (1999) argues, this assumption founders on "the problem of stored beliefs".

15. See Neta (2002; 2003; 2004).

16. See Neta (*forthcoming a*).

Closure: If Stewart has basic knowledge that the table is red, then he would be able to achieve knowledge (by deduction) that the table is not white. So it *is* possible to achieve, by means of deduction, knowledge of a sort that is very similar to what we've been calling "easy knowledge". Indeed, there seem to be clear and unproblematic cases of this. You show me a table that is obviously red. I thereby have basic knowledge that it is red. But from this I deduce, and so come to know, that it's not green. Equally, I can deduce, and so come to know, that it's not white. It seems that I have thereby achieved something like "easy knowledge" that the table is neither green nor white. But if this is so, then why does there seem to be something wrong Stewart's claiming to know that the table is not white with red lights on it?

To answer this question, note that, according to (R), when we bring up the hypothesis that the table is white with red lights shining on it, we move into a context in which we can truthfully claim that Stewart does not know that the table is red (and can truthfully deny that Stewart knows the table to be red).¹⁷ So, if, in this context, we consider Stewart's claim that the table is red, we will naturally (and, in this context, correctly) regard Stewart as making a claim that he doesn't know to be true. That's why, once we consider the hypothesis that the table is white with red lights shining on it, it now seems as if Stewart's inference is epistemically worthless, even though it didn't seem that way before we considered that hypothesis. The hypothesis that the table is white with red lights shining on it is an uneliminated counterpossibility to Stewart's knowing that the table is red, and so it is just the kind of hypothesis the raising of which would tend to generate a context-shift, according to (R). In contrast, the hypothesis that the table is green is not an uneliminated counterpossibility to Stewart's knowing that the table is red. That's why raising the hypothesis that the table is green (for instance) does not also generate a context shift.

Bootstrapping: Recall the example we considered above.

17. This is also predicted by a version of Jonathan Schaffer's contrastivist view. See Schaffer (2004).

- (1a) At t1, it looks to me as if the table is red.
- (1b) At t1, the table is red.
- (2a) At t2, it looks to me as if the cat is on the mat.
- (2b) At t2, the cat is on the mat.
- (3a) At t3, it looks to me as if it's snowing.
- (3b) At t3, it's snowing.
- ...

(conclusion) The way things look to me is (generally) the way they are.

We'll call this inference the "Bootstrapping" inference.

Here again, there seem to be obvious and unproblematic cases of acquiring something very similar to "easy knowledge" by ampliative inference from premises of which we have basic knowledge. Consider the following.

- (1a') At t1, you're saying that the table is red.
- (1b') At t1, the table is red.
- (2a') At t2, you're saying that the cat is on the mat.
- (2b') At t2, the cat is on the mat.
- (3a') At t3, you're saying that it's snowing.
- (3b') At t3, it's snowing.
- ...

(conclusion) You are a reliable reporter of the way things are.

It seems that, by means of this inductive inference, you can move from basic knowledge of all the premises to easy knowledge of the conclusion. But if this is so, then why does there seem to be something wrong with the Bootstrapping inference?

Once again, it helps to appeal to contextualism. According to the contextualist view we've been considering, when we bring up the hypothesis that the way things look to you is not the way they are, we move into a context in which we can truthfully claim that you do not have basic knowledge of the (b) premises of our induction. So, if, in this context, we consider the conclusion of the induction, we will naturally regard you as ignorant of the truth-values of the (b) premises of the induction. That's why, once we consider the conclusion of the

induction, it seems as if you are being epistemically irresponsible in making the induction, even though it didn't seem that way before we considered that conclusion. The conclusion is a hypothesis that, in some sense, you have not ruled out, and so it is just the kind of hypothesis the raising of which would tend to generate a context-shift, according to standard contextualist views. In contrast, the hypothesis that the way you say things are is the way they really are is easily ruled out by the same sorts of considerations that could lead us to believe the (a) and (b) premises of the Bootstrapping induction. That's why considering the latter induction does not generate a context-shift, and the induction seems perfectly in order.

Notice that (R) can help to explain why easy knowledge, as Cohen understands it, strikes us as unattainable while things that are very similar to easy knowledge seem easily attainable. From his basic knowledge that the table is red, Stewart should easily be able to gain knowledge by deduction that the table is not green. By parity, it seems that he should easily be able to gain knowledge by deduction that the table is not white. So what's wrong with his gaining knowledge by deduction that the table is not white with red lights shining on it? The difference here is this: the hypothesis that the table is green is not an uneliminated counterpossibility to S's knowing that the table is red, since the two hypotheses are not introspectively indistinguishable. Neither is the hypothesis that the table is white such an uneliminated counterpossibility, since again, the hypothesis that the table is red and the hypothesis that the table is white are not introspectively indistinguishable. But the hypothesis that the table is white with red lights shining on it *is* an uneliminated counterpossibility to S's knowing that the table is red, so by raising the former hypothesis, we move into a context of appraisal in which we can no longer truthfully affirm, and can truthfully deny, the latter.¹⁸

Again, from our basic knowledge of all of the (b) premises of the testimony induction above, the inductive agent should easily be able to

18. In personal communication, Jonathan Schaffer has suggested that I make this point by saying that each of the hypotheses that we express by saying "the table is red", "the table is green", and "the table is white" includes a tacit domain restriction to the case in which the lighting is non-deceptive. So really the basic knowledge that I'm allowing Stewart can have is knowledge that the table is red (given that the lighting is non-deceptive). Some readers may find this a helpful way of understanding the point I'm making, but I'm reluctant to commit myself to this view about what's included in the hypotheses themselves. For an explanation of my reluctance, see Neta (*forthcoming c*).

gain knowledge by induction that the way you say things are is (usually, at least) the way they are. So what's wrong with her gaining knowledge by induction that the way things look to her is (usually, at least) the way they are? The difference, again, is this: the hypothesis that the way you say things are is not the way they are is not an uneliminated counterpossibility to S's having basic knowledge of the (b) premises. But the hypothesis that the way things look to S is not the way they are is an uneliminated counterpossibility to S's having basic knowledge of the (b) premises. So by raising the second hypothesis (which we do when we consider the conclusion of the ampliative inference), we move into a context in which we cannot truthfully say, and can truthfully deny, that S has knowledge of the (b) premises of the induction.

VII. *Conclusion*

Sometimes, we have basic knowledge. This view has been thought to give rise to intractable problems of easy knowledge. But, as I have argued above, the problems of easy knowledge are not intractable once we accept the independently well-motivated view that ascriptions of evidence (and so of justification and of knowledge) are semantically context-sensitive. We can offer a non-skeptical solution to the problem of easy knowledge without having simply to grant that easy knowledge inferences transmit warrant from their premises to their conclusions, and without having to deny the Closure principles.¹⁹

REFERENCES

- Alston, W. (1983). 'What's Wrong with Immediate Knowledge?', *Synthese* 55, 73–95.
— (1986). 'Epistemic Circularity', *Philosophy and Phenomenological Research* 47, 1–30.
— (1993). *The Reliability of Sense Perception*, Cornell University Press, Ithaca, NY.

19. I am grateful to Michael Bergmann, Martijn Blaauw, Matthew Chrisman, Jonathan Cohen, Adam Leite, Bill Lycan, Eric Marcus, Ted Parent, Jonathan Schaffer, James Shelley, Nicholas Silins, Jay Rosenberg, and especially to Stewart Cohen for helpful discussion.

- Bergmann, M. (2004). 'Epistemic Circularity: Malignant and Benign', *Philosophy and Phenomenological Research*.
- BonJour, L. (1978). 'Can Empirical Knowledge Have a Foundation?', *American Philosophical Quarterly* 15, 1–13.
- Cohen, S. (2002). 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research* 65, 309–29.
- Craig, E. (1990). *Knowledge and the State of Nature*, Clarendon Press, Oxford.
- Davies, M. (2003). 'Solving the Problem of Armchair Knowledge', Hempel Lectures, Princeton University.
- Fogelin, R. (1994). *Pyrrhonian Reflections on Knowledge and Justification*, Princeton University Press, Princeton, NJ.
- Fumerton, R. (1995). *Metaepistemology and Skepticism*, Rowman and Littlefield, Lanham, MD.
- Goldman, A. (1999). 'Internalism Exposed', *Journal of Philosophy* 96, 271–93.
- Harman, G. (1973). *Thought*, Princeton University Press, Princeton, NJ.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford University Press, Oxford.
- Klein, P. (2004). 'Closure Matters: Academic Skepticism and Easy Knowledge', *Philosophical Issues* 14, 165–184.
- Markie, P. (2005). 'Easy Knowledge', *Philosophy and Phenomenological Research*.
- McDowell, J. (1982). 'Criteria, Defeasibility and Knowledge', *Proceedings of the British Academy* 68, 455–79.
- Neta, R. (2002). 'S knows that P', *Noûs* 36, 663–681.
- (2003). 'Contextualism and the Problem of the External World', *Philosophy and Phenomenological Research* 66, 1–31.
- (2004). 'Perceptual Evidence and the New Dogmatism', *Philosophical Studies* 119, 199–214.
- (forthcoming a). 'Epistemology Factualized: New Contractarian Foundations for Epistemology', *Synthese*.
- (forthcoming b). 'Why Moore was not a Fallibilist' in *Themes from G.E. Moore: New Essays in Epistemology and Ethics*, (eds.) Susana I. Nuccetelli and Gary Seay, Oxford University Press, Oxford.
- (forthcoming c). 'Undermining the Case for Contrastivism', *Social Epistemology*.
- Nozick, R. (1981). *Philosophical Explanations*, Belknap Press, Cambridge, MA.
- Pryor, J. (2004). 'What's Wrong with Moore's Argument?', *Philosophical*

- Issues* 14, 349–378.
- Schaffer, J. (2004). ‘From Contextualism to Contrastivism’, *Philosophical Studies* 119, 73–103.
- Sellars, W. (1963). ‘Empiricism and the Philosophy of Mind’ in *Science, Perception and Reality*, Ridgeview, Atascadero, CA.
- Sosa, E. (1997). ‘Reflective Knowledge in the Best Circles’, *Journal of Philosophy* 94, 410–30.
- Van Cleve, J. (2003). ‘Is Knowledge Easy—Or Impossible? Externalism as the Only Alternative to Skepticism’ in *The Sceptics* (ed.) S. Luper, Ashgate, Aldershot.
- Vogel, J. (2000). ‘Reliabilism Levelled’, *Journal of Philosophy* 97, 602–23.
- Williamson, T. (2000). *Knowledge and its Limits*, Oxford University Press, Oxford.
- Zagzebski, L. (1999). ‘What is Knowledge?’, in *The Blackwell Guide to Epistemology*, (eds.) J. Greco and E. Sosa, Blackwell, Malden, MA and Oxford.