

Why Basic Knowledge is Easy Knowledge

STEWART COHEN
Arizona State University

The problem of easy knowledge arises for theories that have what I call a “basic knowledge structure”. S has basic knowledge of P just in case S knows P prior to knowing that the cognitive source of S’s knowing P is reliable.¹ Our knowledge has a basic knowledge structure (BKS) just in case we have basic knowledge and we come to know our faculties are reliable on the basis of our basic knowledge. The problem I raised in “Basic Knowledge and the Problem of Easy Knowledge”² (BKEK) is that once we allow for basic knowledge, we can come to know our faculties are reliable in ways that intuitively are too easy. This raises a serious doubt about whether we had the basic knowledge in the first place.

In “Easy Knowledge”, Peter Markie argues that BKS theories do not face any problem concerning easy knowledge.³ I argued that the problem arises in two forms, and Markie takes issue with both. I will argue that Markie’s defense of BKS theories fails.

Easy Knowledge: Deductive Closure

The first version of the easy knowledge problem arises in connection with a very plausible version of the deductive closure principle for knowledge.

¹ I originally defined “basic knowledge” in this way to avoid the problem of the criterion. As Nico Silins has pointed out to me, this would allow for a BKS view that requires that one have evidence (short of knowledge conferring) for the reliability of one’s knowledge source prior to having knowledge by that source. For this view to get anywhere, it would have to be possible to get knowledge of the reliability of a source without having knowledge by that source. Otherwise, a regress would ensue. As near as I can tell, a view of this kind would still face the easy knowledge problem, so long as the prior evidence for reliability is less than required for knowledge of reliability. Also, in my original definition of “basic knowledge”, I referred to the source of S’s *belief* in P, rather than, as I now do, the source of S’s knowledge of P. This change is necessary to avoid counting as basic, knowledge I acquire from a source that is distinct from the belief source. Thanks again to Nico Silins.

² Cohen (2002)

³ Markie (2005)

DC: If S knows P and S competently deduces Q from P, then S knows Q.⁴

To illustrate the problem, I consider a case where my son wants a red table for his room, but is worried that the table that looks red might be white with red lights shining on it. I try to reassure him by the following reasoning.

- (I) (1) The table looks red
- (2) The table is red
- (3) If the table is red, then it is not white with red lights shining on it
- (4) The table is not white with red lights shining on it.

On standard BKS theories, I can come to know (2) on the basis of (1). Since I know (3) *a priori*, then given DC, I can come to know (4), on the basis of (2) and (3). I claim that (I) is clearly an inadequate response to my son. The problem for BKS theories is that given the deductive closure principle, the theory has to allow that my reasoning is acceptable.

Markie argues that this case does not show there is a problem for BKS theories. For there is an alternative explanation for our judgment that my response to my son is inadequate:

Worried about deception, the son wants to believe that the table is not white and illuminated by red lights in order to believe it is red. In (I), the father offers him as a reason the very point he is questioning, that the table is red. He begs the question (p. 409).

Markie raises an interesting point. As he notes, one can acquire knowledge of P by reasoning in a certain way, even though in certain contexts, reasoning this way in defense of P would be dialectically ineffective. Surely this is how, in general, we must view anti-skeptical reasoning. Even if such reasoning fails to convince a stubborn skeptic, we can still come to know on the basis of that reasoning. Markie thus seeks to explain our intuition that the father's response is inadequate to address the son's concern, while still allowing that the father can know by such reasoning.

Though Markie's strategy may partly explain our response to the reasoning in (I), I do not think it can fully explain it. Markie argues that the father begs the question against the son by assuming a premise the son rejects, viz., that the table is red. But as I have constructed the case, the father does not *assume* that the table is red. Rather he infers that the table is red from the fact

⁴ This version of the closure principle is due to Timothy Williamson (2000). It is roughly equivalent to the version I proposed in Cohen (2002), but I think more perspicuous.

that it looks red.⁵ So the father begs the question against the son by reasoning in a way that the son rejects. But even if anti-skeptical reasoning is dialectically ineffective against a skeptic, if it is to have any anti-skeptical force *for us*, it must strike *us* as good reasoning. And insofar as it does, we must view that skeptic as wrong for not accepting that reasoning. But is this our attitude toward the reasoning in (I)? Do we think that my son is wrong for not accepting that reasoning? Speaking for myself, rather than thinking that my son is being irrational for not being convinced by the reasoning, I find myself in sympathy with him. The reasoning is not cogent *for me*.

Consider how a case would look where our only reservation about it concerned its lack of dialectical cogency: Suppose my son is worried that our car is going to run out of gas. He asks me if I know whether the gas tank is full. I say, "Yes it is—the fuel gauge says it's full and the gauge is very reliable". My son replies, "Why should I believe the tank is full, just because the reliable gauge says so?" I reply, "Look—the gauge says it's full and the gauge is reliable. So you should believe the tank is full". Here, I have clearly begged the question against my son. He questions whether I can infer that the tank is full from the fact that the reliable gauge says it is full. And I appeal to precisely that inference in response. But does my response to my son strike us as suspect? On the contrary, my son appears to us to be irrational in refusing to accept it. This remains true despite the fact that the reasoning begs the question against him. This shows that we can still judge that a piece of reasoning is good, even when it begs the question in a particular dialectical context.

Now consider again the case involving the reasoning in (I). In just the same way, the father begs the question against the son by appealing to the very reasoning he questions. But as I noted, in this case, it does not appear to us (or to me at least) that the son is being irrational in rejecting the reasoning. This indicates that we find the initial reasoning to be suspect.

Consider what happens if we remove my son from the case. Suppose I am interested in buying a red table and I am anxious to avoid buying a white table that looks red (in the showroom) because it is illuminated by red lights. In order to convince myself that a particular table is not deceptively illuminated, I appeal to the reasoning in (I): "Hmmm, the table looks red, so it is red, so it is not white with red lights shining on it." Suppose further that as a result of going through this reasoning, I become convinced that the table is not white with red lights shining on it. Presumably, the intuition remains that there is something wrong with my having reasoned in this way. Can we appeal to the dialectical context to explain why? Of course here there is only one reasoner. I suppose we could view this as a case where I am having a

⁵ I don't mean to be arguing that in color perception, we always engage in explicit inferences of this sort

discussion with myself. Perhaps then we could say that by appealing to the reasoning in (I), I have begged the question against my former skeptical self. But as Markie notes, it does not follow from this that my current self cannot come to know by this reasoning. So, from the perspective of my current self, if BKS theories are correct, my reasoning is impeccable. But again, surely there is something unsatisfactory about the way I have convinced myself that the table is not deceptively illuminated. Again, the appeal to dialectical context cannot explain our intuitions.

Finally we could redescribe the case to eliminate entirely any dialectical context. Suppose my son is not worried about the possibility that the table is deceptively illuminated. He accepts that I know that the table is red and not white with red lights shining on it. He is just curious about *how* I know it. I respond in just the same way. "Oh, that's easy. It looks red, so it is red, so it is not white with red lights shining on it." By my lights, this does not change the intuition that there is something unsatisfactory about my response. Of course, it could be that *we* find the response unsatisfactory simply because *we* are now questioning whether the father can, in this way, know the table is red. But that is just to say that the case makes *us* doubt the truth of BKS theories. And that is the whole point of my presenting the case.

My second reason for denying that one could, in this way, come to know that the table is not white but illuminated by red lights is that it would make our inductive evidence concerning the likelihood of such a deception irrelevant to our knowing we are not so deceived. Markie responds that, on a BKS view, one's inductive evidence is still necessary for formulating non-question-begging responses to a skeptic who questions the reliability of color vision.

But even if this is true, the fact remains that on a BKS view, our inductive evidence against the likelihood of encountering a white table illuminated by red lights is irrelevant to our coming to know that a table that looks red is not in fact a white table illuminated by red lights.⁶ And surely this is a strange result.

Easy Knowledge Without Basic Knowledge?

Markie also argues that his proposal for handling my case can handle another version of the easy knowledge problem that arises even for theories that deny basic knowledge. Suppose I know on the basis of substantial evidence that my color vision is reliable. Again my son wonders whether the red-appearing table might be white with red lights shining on it. I respond with the following reasoning;

(II) (1) [My inductive evidence for the reliability of color vision]

⁶ I say irrelevant to our *coming to* know since on the BKS view, our inductive evidence could supply us with an additional source of knowledge.

- (2) Color vision of the sort I'm employing right now is reliable and the table looks red.⁷
- (3) The table is red
- (4) If the table is red, it is not white with red lights shining on it.
- (5) The table is not white with red lights shining on it.

Markie argues that "there are certainly ways in which the father can use his inductive evidence to provide a satisfying response to his son's concern, but (II) is not one of them (p. 410)." This is the case, even though my knowledge that the table is red is based on the reliability of color vision.

Before we consider how to respond to this case, we need to be clear about the problem Markie thinks is posed by the reasoning in (II). At first blush, steps (2), (3) and (4) appear to be superfluous. If (1) supports (2), then my inductive evidence will provide a basis for thinking that deceptive lighting is very improbable. But then (1) will support (5) directly.⁸

Markie says the reasoning is an inadequate response because my son should reply,

I know that your color vision is reliable, but I'm concerned that in this particular case, it may be inaccurate. It is not an adequate answer to my concern to tell me that my worry does not apply because...you offer me...inductive evidence for perception's general reliability. What I am questioning is whether despite all that, the table may still, in this case, be white and illuminated by red lights (p. 410).

So my son's concern is that even though my inductive evidence makes the probability of deceptive lighting very low, I nonetheless fail to know that *in this very case*, the improbable deception is not occurring. My son is questioning whether (1), even though it supports (the first conjunct of) (2), really does support (5). I agree that (II) appears to be an inadequate response to my son's query, understood in the way Markie indicates. The problem is that even if (1) does not by itself support (5), if (1) supports (2), then the reasoning can proceed all the way to (5) via (2)-(4).

Markie points out that my appeal to Sosa's distinction between animal and reflective knowledge cannot handle this case⁹. For I do have reflective knowledge that the table is red and so deductive closure applies.

⁷ Markie leaves out the second conjunct but presumably it is required.

⁸ The idea here is that if my inductive evidence supports the reliability of color perception, then it supports that deceptive lighting is very unusual. And that supports that I am not now seeing a deceptively lit table. Perhaps it would be clearer if I said that (1) along with the second conjunct of (2) supports (5).

⁹ Sosa (1997).

Markie then notes that his treatment of my original case will handle his new case as well. We can say that the reasoning in (II) does give me knowledge, while explaining our intuition that the reasoning is inadequate, by noting that it begs the question against my son. So on Markie's view, even on the assumption that (1) does not by itself support (5), the reasoning via (2)-(4) can in fact give me knowledge of (5). The only problem with the reasoning is its unsuitability for the dialectical context.

I have argued that this kind of response to my original case fails and the same considerations apply here. The reasoning seems unsatisfying to us, not merely inadequate dialectically against my son. And we could redescribe the case so I am questioning my own reasoning.

Contrary to Markie's view, I hold that if (1) is not by itself sufficient for me to know (5), then reasoning via (2)-(4) cannot give me knowledge of (5) either. So what should we say about (II)? I have argued that the easy knowledge problem appears to arise from our rejection of principle KR.¹⁰

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

But now we have an argument that even if we accept KR, we still face the problem of easy knowledge. According to this argument, evidence can give us knowledge that perception is reliable, without giving us knowledge that perception is veridical on a particular occasion (provided the reliability is not perfect). For example, I could know that color vision is in general reliable without knowing that on a particular occasion when a table looks red, the table is not deceptively illuminated. But if knowing my color vision is generally reliable allows me to know that the red-appearing table is red, then there is nothing to stop me from employing the reasoning in (II) to gain (easy) knowledge that the table is not deceptively illuminated on this occasion. And so, even if we accept KR, we do not avoid the problem of easy knowledge.

This kind of skeptical challenge is quite general. Virtually anything we know on the basis of evidence, is rendered at best highly probable on that evidence. The question raised by Markie's case is this: How do we know that in a particular case where our evidence makes P highly probable, that some improbable alternative to P does not obtain?

Markie is correct to point out that evidence for the general reliability of color perception does not *seem* to give us knowledge that color perception is not deceptive on some particular occasion. And I have argued that Markie's dialectical explanation cannot explain this intuition away.

¹⁰ Cohen (2002).

Where does this leave us? I think that when we know P is highly probable, we *can* know that some highly improbable alternative to P does not obtain. We can know it precisely because we know that it is highly improbable. So the evidence for the general reliability of color perception does enable us to know that color vision is not deceptive on this particular occasion. But then how do we explain why it seems reasonable for my son to question this? More generally, how do we explain the intuition that evidence that makes P probable is not sufficient to know that on this particular occasion, the improbable (some sub-case of not-p) is not occurring?

Here an appeal to a contextualist account of knowledge ascriptions is useful. According to Contextualism, the truth-value of knowledge ascriptions depends on context-sensitive standards. On my version, the standards govern how strong one's evidence must be in order for one to know.¹¹ One feature of context that can affect the standards is the salience of error possibilities. Thus "S knows the table is red" can be true at an everyday context and yet false at a higher-standards context where certain error possibilities are salient. If this is correct, then we can explain away the intuition that we cannot, on the basis of knowing that perception is reliable, know that perception is not deceptive on some particular occasion. Consider everyday contexts where error possibilities are not salient. At such a context, the standards allow that our inductive evidence makes it sufficiently probable that color perception is reliable for S to know that the table is red on the basis of its looking red. But then at those same contexts, our inductive evidence for the reliability of color perception is sufficient for us to know that the table is not white and illuminated by red lights, i.e., that color vision is not deceptive on this occasion

Why then do we have the intuition that the evidence for the reliability of color perception is not sufficient for us to know that on this particular occasion, the lighting is not deceptive? Once the possibility that color perception is deceptive has been made salient, we shift to a stricter context where our evidence for the general reliability of perception is *not* sufficient for us to know that on this particular occasion the table is not white with red lights shining on it. So our intuition that we do not know that the lighting is not deceptive on this occasion is explained by our having shifted to a high standards context where we do not in fact know that the lighting is not deceptive. Moreover, at this higher standards context, we do not know that the table is red on the basis of our inductive evidence concerning the reliability of color perception. So according to contextualism, in everyday contexts when error possibilities are not salient, our evidence for the reliability of color perception is sufficient for us to know both that the red-looking table is red and that the red-looking table is not white with red lights shining on it. But at higher

¹¹ Cohen (1988), (1999). For other versions of Contextualism, see DeRose (1995) and Lewis (1996).

standards context where error possibilities are salient, our evidence is not sufficient for us to know either of these propositions.

But if Contextualism can rescue a non-BKS account of our evidence from Markie's problem case (II), why can't Contextualism rescue the BKS account from my original case (I)? Why not say that the reason we have the intuition that my response to my son is inadequate is that once my son raises the possibility of deception, there is a shift to a high standards context where the table's looking red is not sufficient to know it is red? This would allow that the table's looking red is sufficient for knowing that it is red at everyday contexts.

In BKEK, I noted that a BKS theorist might be tempted to give precisely this kind of contextualist response to my original case. And let me state for the record that I would welcome this application of Contextualism, if it could be sustained. I rejected a contextualist treatment of the case on the grounds that, by my lights anyway, the BKS account of how we know that the table is not deceptively illuminated is implausible, even when the context is held fixed.¹² But then why can I appeal to Contextualism to handle Markie's case of easy knowledge that arises even for non-BKS theories?

Of course this is a general problem for contextualist treatments of skeptical problems. Though Contextualism enables us to explain away certain skeptical intuitions by appealing to contextual shifts, the question can remain whether the original intuition depends on such a shift. So let me try to provide some independent considerations in support of our skeptical intuitions about my original case.

Basic Knowledge and Evidential Probability

Let us return to the case of the red table. We can stipulate that before I look at the table, I have no idea what color it is. Moreover on a BKS account, we can stipulate that before I look at the table, I do not know whether it is more likely that a red appearing table is red, or non-red but deceptively illuminated by red lights. When I actually look at the table, I acquire the evidence that the table looks red. On the BKS account, I can, in virtue of acquiring this new evidence, come to know that the table is not non-red but deceptively illuminated by red lights. But though this new evidence counts in favor of the table's being red, it is not clear why it does not also count in favor of the table's being non-red but deceptively illuminated to appear red. After all, if one were looking for such a table, one would eliminate the ones that do not look red and investigate further the ones that do look red. But if the evidence that the table looks red counts in favor of its being non-red but illuminated to appear red, then it is puzzling how acquiring this evidence could enable me to know that the table is *not* non-red but illuminated to appear red.

¹² See endnote 15 of Cohen (2002).

We can express this point more precisely by considering the evidential probabilities from a BKS perspective.¹³ We can stipulate that before I look at the table, my evidential probability for the table's appearing red is .5. So my evidential probability for the table's not appearing red is .5 as well. And consistent with our previous stipulation, we can say that before I look at the table, my conditional probability for the table's being red, given that it appears red, is .5. And so my conditional probability for the table's being non-red but red-appearing, given that it appears red, will be .5 as well. The prior probability for the table's being non-red but red-appearing will be the product of Pr (the table appears red) and Pr (the table is non-red but red-appearing/the table appears red). That product is .25. When I do look at the table, we can suppose that my probability for the table's looking red goes to 1 (or negligibly close to it). Then the probability that the table is non-red but red-appearing (after it appears red to me) will be equal to the prior conditional probability of the table's being non-red but red-appearing, given that it appears red. That probability we said is .5.

But then this gives us a very strange result: Before I look at the table, I do not know that it is not non-red but red-appearing. (After all, the probability that this is the case will be only .75, given our stipulations.) But we have seen that the new evidence I get when I look at the table viz., that it is red-appearing, actually raises the probability that the table is non-red but red-appearing from .25 to .5. Yet according to BKS theories, on the basis of this new evidence, I can come to know that the table is not non-red but red-appearing (by inferring it from the fact that the table is red). So according to BKS theories, acquiring evidence which raises the probability that the table is non-red but red-appearing provides a basis for my coming to know that the table is *not* non-red but red-appearing. Surely any conception of evidence that generates this result is untenable.

This problem for BKS theories clearly indicates that an application of the contextualist strategy for explaining away skeptical intuitions is not called for in my original case (I). But I see no barrier to applying that strategy in case (II).

Easy Knowledge: Bootstrapping

The bootstrapping problem was first raised by Jonathan Vogel and Richard Fumerton in connection with Reliabilism.¹⁴ In BKEK, I argued that the problem could be generalized to evidentialist theories that allow for basic knowledge. I first presented a case that demonstrates the problem for Reliabilism:

¹³ I am indebted here to conversations with John Hawthorne and Frank Arntzenius. Hawthorne makes essentially the same point in Hawthorne (2004) p. 73-5.

¹⁴ Vogel (2000) and Fumerton (1995).

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

I also argued that if we modify the case so I reason, "That screen is red and it looks red", then the case shows that evidentialist BKS theories are committed to the same result.

Markie's first concern about this kind of case is that it may be impossible. As I present the case, I make color judgments without having any evidence for the reliability of color perception. But according to Markie, in order to be able to make color judgments, one must have learned how to identify colored objects through perception. And learning how to identify colored objects through perception will involve getting some evidence for the reliability of color perception.

I am going to avoid the thorny issue about what is necessary for concept acquisition and grant for the sake of argument that Markie is correct. But even if one must have evidence for the reliability of color perception in order to learn how to identify colors, the bootstrapping problem remains. For we could stipulate that my son and I have forgotten whatever evidence we once had, and so are unsure whether our color vision is reliable. Thus my son could raise the question of whether color vision is reliable and I could answer him by going through the slide-show procedure.

One might object that one must retain one's evidence for the reliability of color perception in order to retain the concept of red. While this strikes me as implausible, I could grant the point and still raise the bootstrapping problem. As Markie himself notes (see note 6), there is nothing in the BKS account that would require us to *base* our color judgments on whatever evidence we have for the reliability of our color perception. Suppose then I run through the slide show making color judgments without basing them on any of my evidence of my reliability. The untoward consequence remains—according to the BKS account, I would thereby generate evidence for the reliability of my color perception.

Next, Markie argues that my own solution to the bootstrapping problem fails along with another solution he considers. He then concludes that

...the best strategy is to accept that, through basic knowledge we can use a faculty to gain evidence of its own reliability. Assuming that Cohen's example is possible, the father can use his color perception to gain evidence of the reliability of color perception (p. 415).

Markie here means to be allowing that the bootstrapping procedure I go through in the slide show *does* generate evidence for the reliability of perception. His explanation for why it strikes us as illegitimate is the same for the bootstrapping case as it was for the closure case. Because my son is questioning the reliability of color perception, the justification I have for believing my color perception is reliable cannot be offered to my son without begging the question against him.

In response, I should first note that I do not deny that one can use perception to gain evidence for the reliability of perception. Clearly we do know that perception is generally reliable. And as I noted in BKEK, short of holding that we know *a priori* that perception is reliable, there is no way of denying that we come to know that perception is reliable by in part relying on perception. In the end, I defend a view according to which we do come to know perception is reliable by relying, in part, on the deliverances of perception. The bootstrapping problem for BKS theories is *not* that they allow us to know perception is reliable by relying on perception. Rather the problem is that BKS theories allow us to rely on perception in an intuitively illegitimate way, viz., by bootstrapping.

Can we explain away the intuitive illegitimacy of bootstrapping by appealing to dialectical concerns about question-begging? As I argued in the case of the reasoning in (I) and (II), I don't think that the dialectical inadequacy against my son can fully explain our intuition that the reasoning is objectionable. For the reasoning is not cogent *for us*.

Moreover, as was true for the deductive closure case, there is nothing in the bootstrapping example that requires that my son is questioning the reliability of color perception. We can stipulate that he believes that I know color perception is reliable. He simply wants to know how I know it. Let us suppose that I believe that color perception is reliable on the basis of the procedure I performed in the slide show example. Nonetheless, it seems that if I run through the slide show again to demonstrate how I know color perception is reliable, my son should not be satisfied with my response. He should think that I'm merely spinning my wheels. But I have not begged any question against him.

Solutions and Problems

So if the problem of easy knowledge remains, what is the solution? In BKEK, I distinguished, following Ernest Sosa, between reflective knowledge and animal knowledge. I argued that though on the one hand, we must deny that we can have basic knowledge in order to avoid easy knowledge, on the

other hand it seems hard to deny that children have knowledge that appears to be basic. My strategy was to deny that one can gain *reflective* basic knowledge while allowing that one can acquire *animal* basic knowledge. Then I suggested that animal knowledge cannot combine individually with self-knowledge (in this case, the knowledge that the slide looks red to me) to generate inferences.

Markie argues, I believe correctly, that this solution is problematic. He gives a clear example where animal knowledge combines with knowledge of preferences to generate an inference that looks perfectly acceptable. Thus Markie shows that I must modify my account.

When discussing a possible solution to the easy knowledge problem, Markie considers the following principle

IP: Where a belief gains it *prima facie* justification for S just from that 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 (p. 414).

This principle would block the bootstrapping reasoning in the slide show example. Markie rejects this principle on the plausible grounds that if a belief B is an instance of knowledge (or reasonable belief), then the evidential value of B should not be affected by its source. "Just as money, however gained, still spends the same, so too reasonable beliefs, however gained still epistemically support the same beliefs (p. 415)." Moreover, if we accept IP, it is hard to see how we could ever come to know that perception is reliable.

But suppose we interpret IP as a principle governing animal knowledge/justification. The upshot would be that one cannot appeal to animal knowledge produced from a particular faculty to gain knowledge of the reliability of that faculty. Since at best, one can have animal knowledge that the slide is red, simply on the basis of its looking red, this would block the kind of bootstrapping reasoning that we find to be illegitimate.

Do the objections to IP apply when it is interpreted as a principle governing animal knowledge? There is no problem with our gaining knowledge of the reliability of perception since IP would allow us to appeal to reflective perceptual knowledge. What about Markie's worry that reasonable beliefs however gained still epistemically support the same beliefs. I concede that this point is intuitively very plausible. But once we distinguish between animal knowledge/reasonable belief and reflective knowledge/reasonable belief, we can interpret that intuitive plausibility as pertaining to reflective knowledge/reasonable belief.¹⁵ This allows us to claim that IP holds for animal knowledge.

¹⁵ This is the line I took in Cohen (2000) regarding animal knowledge and the closure version of easy knowledge. I have to take a similar line concerning animal knowledge and

I do not deny that this proposal is *ad hoc*. But at present, it is the best I can do.¹⁶

the rule of conditionalization since as Nico Silins has pointed out to me, the problem concerning conditionalization that I raised for basic knowledge applies to animal knowledge as well.

¹⁶ I thank Frank Arntzenius, Tom Blackson, John Devlin, Richard Fumerton, John Hawthorne, Peter Markie, Ram Neta, Jonathan Schaffer, Nico Silins, and Roger White, for helpful comments.

References

- Cohen, Stewart (1988) "How to be a Fallibilist" in *Philosophical Perspectives*, 2, Tomberlin, James (ed.)
- (1999) "Contextualism, Skepticism, and the Structure of Reasons" in *Philosophical Perspectives*, 13, Tomberlin, James (ed.)
- (2002) "Basic Knowledge and the Problem of Easy Knowledge" *Philosophy and Phenomenological Research*, 65
- DeRose, Keith (1995) "Solving the Skeptical Problem" *Philosophical Review*, 104
- Fumerton, Richard (1995) *Metaepistemology and Skepticism* (Lanham, MD: Rowman and Littlefield)
- Hawthorne, John (2004) *Knowledge and Lotteries* (Oxford: Oxford University Press)
- Lewis, David (1996) "Elusive Knowledge" *Australasian Journal of Philosophy*, 74
- Markie, Peter (2005) "Easy Knowledge" *Philosophy and Phenomenological Research*, 70
- Sosa, Ernest (1997) "Reflective Knowledge in the Best Circles." *Journal of Philosophy*, 94
- Vogel, Jonathan (2000) "Reliabilism Leveled" *Journal of Philosophy*, 97
- Williamson, Timothy (2000) *Knowledge and its Limits* (Oxford: Oxford University Press)