Probability and Scepticism

Brian Weatherson *

May 25, 2009

1 The Humean Argument

I've been interested in recent papers in the following argument, derived from Hume. In the argument E is the agent's evidence, and H is some hypothesis derived by ampliative reasoning from her evidence. [1]

1. It is not possible for the agent to know $E \supset H$ *a priori*.
2. It is not possible for the agent to know $E \supset H$ *a posteriori*.
3. So, it is not possible for the agent to know $E \supset H$

If we add as an extra premise that the agent does know H, then it is possible for her to know $E \supset H$, we get the conclusion that the agent does not really know H. But even without that closure premise, or something like it, the conclusion seems quite dramatic.

One possible response to the argument, floated by both Descartes and Hume, is to accept the conclusion and embrace scepticism. We cannot know anything that goes beyond our evidence, so we do not know very much at all. This is a remarkably sceptical conclusion, so we should resist it if at all possible.

A more modern response, associated with externalists like John McDowell and Timothy Williamson, is to accept the conclusion but deny it is as sceptical as it first appears. The Humean argument, even if it works, only shows that our evidence and our knowledge are more closely linked than we might have thought. Perhaps that's true because we have a lot of evidence, not because we have very little knowledge.

There's something right about this response I think. We have more evidence than Descartes or Hume thought we had. But I think we still need the idea of ampliative knowledge. It stretches the concept of evidence to breaking point to suggest that all of our knowledge, including knowledge about the future, is part of our evidence. So the conclusion really is unacceptable. Or, at least, I think we should try to see what an epistemology that rejects the conclusion looks like.

2 Probability and Premise 2

Rejecting the conclusion would be easy if it was easy to reject the premises. But in fact there are quite strong defences of each of the premises. Let's look at some of them.

*Thanks to David Chalmers and Crispin Wright. References etc incomplete.

[1] On how closely this argument resembles Hume's argument for inductive scepticism, see two recent papers by Samir Okasha, "Does Hume's argument against induction rest on a quantifier-shift fallacy?" and "What did Hume Really Show about Induction?". I'm going to assume throughout that we aren't dealing with the special case where the prior credence of E is 0, or of H is 1. That will do some work in section 2.

The simplest argument in favour of premise 1 uses a little bit of empiricism. It could turn out to be true that $E \supset H$. What could turn out to be true can only be known *a posteriori*. So we can't know *a priori* that $E \supset H$. The crucial premise there, about the limits of the *a priori*, is the distinctively empiricist assumption, but it is shared by a lot of contemporary philosophers.

The simplest argument in favour of premise 2 uses a little bit of rationalism, though I think it takes a little more to see that it is a rationalist assumption. Here's the argument in premise-conclusion form; we'll go through each of the premises at some length below. So as to avoid confusion with the Humean argument, I've named the premises rather than numbered them.

Updating is Conditionalisation (UIC) If we use Cr to measure our rational agent's credences, and Cr_Y to be her credences after updating with evidence Y, then $Cr_Y(X) = Cr(X|Y)$ for all X, Y.

Learning Doesn't Lower Credence (LDLC) It is impossible for Y to be the basis for learning X, i.e., coming to know that X is true, if $Cr_Y(X) < Cr(X)$.

Credences are Classical Probabilities (CCP) Cr is a classical probability function, and it is a theorem of the classical probability calculus that $Pr(E \supset H|E) < Pr(E \supset H)$ when $Pr(E) > 0$ and $Pr(H) < 1$.

Humean Conclusion (HC) So, it is impossible to learn $E \supset H$ on the basis of evidence E.

I'm assuming here that if E is our total evidence, and we know something *a posteriori*, then we know it on the basis of E. (There's probably a strict sense of 'know on the basis of' according to which we don't know H on the basis of E, but on the basis of some subset of it. I'm using the term more liberally here; our knowledge is based in E if E includes all of our grounds for that knowledge, even if E includes much else besides.) So (HC) entails premise 2 of the Humean argument directly. And (HC) follows from (UIC), (LDLC) and (CCP). So if someone wants to reject the Humean argument at premise 2, they better reject one of these three principles.

None of the three is beyond dispute. Any argument against using classical logic when reasoning about fundamental matters will be a reason for rejecting (CCP). David Jehle and I show in some forthcoming work that $Pr(E \supset H|E) < Pr(E \supset H)$ fails for some intuitionist probability functions, so that's one possible way out of the argument. But I imagine it will be too costly a way out for various philosophers to give up classical logic just to escape this argument.

In "The Bayesian and the Dogmatist", I argue that philosophers sympathetic to empiricism (broadly construed) should reject (UIC). That's because (UIC) embodies a very implausible picture of the relationship between evidence and hypotheses. In particular, it assumes that our prior credences can contain all the information we need about how evidence supports hypotheses. When we get evidence, we just react in a way that was antecedently deducible. That seems implausible to me. We can't always know *a priori* what the right way to respond to certain evidence is. And that implies we shouldn't always update by conditionalisation.

But rejecting (UIC) comes at a cost. I also argue in "The Bayesian and the Dogmatist" that when E does not ground new knowledge about fundamental relationships between evidence and justification, then we should update by conditionalisation. So if (UIC) fails, we have to give up the idea that we can know about the fundamental epistemological relationships *a priori*; rather, we have to learn them. And that's true, I claim, even if those relationships hold of necessity. Put another way, if we accept a rationalist principle about fundamental epistemology, namely that we can know its most important contours *a priori*, then (UIC) arguably follows, and with it (HC), and perhaps with it the full Humean sceptical argument.

I think this poses a deep dilemma for certain philosophers, in particular philosophers who want to tie *a priority* too closely to modality. If we can only know contingent truths *a priori*, then premise 1 of the Humean argument is true. If we can know necessary truths *a priori*, and the fundamental epistemological relationships are grounded in necessary truths, then we can know those relationships *a priori*, and hence (UIC) holds, and hence premise 2 of the Humean argument holds. But now disaster strikes, for the Humean argument is complete. We must, if we are to break free of the Humean argument, break the connection between *a priority* and modality. So there are deep issues at play here.

But wait! There was one other assumption that we might reject. The argument for premise 2 also uses (LDLC). If we can overturn (LDLC), then we can resist the argument for premise 2 without rejecting (UIC), and hence without being led to the hard question of what could justify the failure of (UIC). That's the position I want to investigate here.

The plan of attack is as follows. I'll discuss a stylised example where (LDLC) seems to fail. It seems to me that if (LDLC) fails at all, it should fail in cases like this one. In any case, I hope the response I'm going to make to this example generalises to all putative counterexamples. That response will turn on using an interest-relative theory of knowledge. I'll sketch, quite briefly, some motivation for such a theory, and my preferred version of that theory. The theory will be incomplete. I'd like to say that's because of space considerations, but really it's because my understanding is incomplete. I think I can say enough about the theory though to say why we should reject the counterexamples to (LDLC). And that's all I need here to motivate (LDLC), and with it the need to say something about where the Humean argument fails.

3 Knowledge and Probability

David Chalmers and Crispin Wright both suggested to me that (LDLC) could fail in the kind of case I'll describe in this section. The case concerns testimony from a source not certain to be reliable or knowledgeable, and we need a way to model that. I'll assume that if Ra is the proposition that R is a knowledgeable testifier, and Sap the proposition that a said that p, then our agent's credences satisfy the following constraints for any testifier a.

- $Cr(p|Ra \wedge Sap) = 1$
- $Cr(p|\neg Ra \wedge Sap) = Cr(p|\neg Sap)$

That is, testimony from a knowledgeable source is maximally valuable testimony, while testimony from other sources has no evidential value. Both of these assumptions are a little extreme, but more moderate models will also generate the kind of example we're interested in here. [2]

The case concerns a lottery that is based around a series of coin flips. Each lottery ticket consists of a 20-character string of H's and T's. A fair coin is flipped 20 times in a row. If the sequence of H's and T's on one's ticket matches the sequence of Heads and Tails that come up as the coin is flipped. The rational agent

[2] And I think this kind of model is more realistic than a model that is based around Jeffrey-conditionalising, where we have to specify in advance what the posterior probability of some salient proposition is. That's not required here; the posterior probability of p is an output of the prior probabilities of p and Ra, not an input to a Jeffrey-conditionalising formula.

has one ticket in this lottery, so their initial credence that they will lose the lottery is $1 - 2^{-20}$. Let X be the proposition that they will lose the lottery.

The agent will get some testimony from two sources, first b, then c. The agent's prior credence in Rb is 1. That is, she takes b to be perfectly reliable. And her credence in c is 0.99, which is reasonably high. Indeed, it is arguably higher than most everyday testifiers. But she does allow there is a non-zero probability that c's vision was inaccurate, or that their memory was inaccurate, or that they are being deliberately misleading, or that any one of the myriad ways in which individual testifiers fail to be accurate infected c's testimony. The agent then gets the following two pieces of evidence.

- The agent is told by b that the first 19 characters on their ticket match the first 19 flips of the coin.
- The agent is told by c that the last character on their ticket does not match the last flip of the coin.

In both cases we'll assume that the testifiers know the truth of their assertions, though we won't make any assumptions yet about whether the agent shares in this knowledge. After she gets the first piece of evidence, her credence in X drops to 0.5. After she gets the second piece of evidence, her credence in X rises back up to 0.995. That's high, but notably it is less than her prior credence in X.

Still, we might think that the agent is now in a position to know X, and she wasn't before getting this evidence. She has learned that her ticket lost from a knowledgeable source. (Strictly, she has learned something that entailed this, but this doesn't affect the overall argument.) To be sure, she has some minor reservations about the reliability of this source, but those reservations are no greater than most of us have about the testimony we get from friends and acquaintances everyday. And we typically take that testimony to produce knowledge. So it looks like, if Y is the combination of these two pieces of testimony, then Y lowers her credence in X, as we'll put it, it makes X less credible, but it also grounds knowledge of X. That's a counterexample to (LDLC), or so it looks.

Someone might object here that for many everyday pieces of knowledge, the prior credibility of our testifier is greater than 0.99. That doesn't mean the testifier is right 99% of the time, just that on this occasion the credibility of their knowledgeability is greater than 0.99. I'm sympathetic to this line of criticism - I think we often overestimate the likelihood of error in everyday settings. But I don't think it matters much here. For one thing, we often learn things by testimony whose prior credence is much lower than 0.99. For another, we could make the prior credence in c's knowledgeability as high as $1 - 2^{-19}$ without affecting the argument. And that's a very high degree of credibility indeed. It seems to me that c is a lot like an ordinary testifier, and rejecting c's testimony as a grounds for knowledge puts one at grave risk of embracing an overarching scepticism about testimonial knowledge. That is a sufficient reason to stay away from this kind of objection.

It might also be objected that what we have here is a case where there is no single piece of evidence that both lowers the credibility of X and grounds knowledge of X. True, if we take Y to be the combination of the two pieces of evidence the agent gets, then Y both lowers the credibility of X and grounds knowledge of X. But that's because Y has two parts, and one part lowers the credibility of X while not grounding knowledge of it, and the other raises the credibility of X and grounds knowledge of it. If we restrict our attention to single pieces of evidence, says the objector, then (LDLC) is clearly true, and is untouched by this objection.

I think this objection is correct as far as it goes. It will be tricky to fill in the details of the objection, since it will be tricky in all cases to individuate pieces of evidence. (I like a causal theory of evidence on which individuating evidence isn't as hard as it looks on some other theories, but that's a topic for another day.) But the problem with this objection isn't that it relies on a tricky individuation of pieces of evidence. Rather, it is that the restricted version of (LDLC) that the objector falls back on isn't sufficient to save the Humean argument for scepticism. The Humean argument, which is our focus in this paper, asks whether our total evidence E could ground knowledge in propositions like $E \supset H$. And we're imagining a proponent of the Humean argument who answers "No" by appeal to (LDLC). Since our total world evidence is not a single piece of evidence, for this response to work, we need a version of (LDLC) that applies to compound evidence, such as the testimony that the agent receives from b and c.

So if we are to defend (LDLC), and hence defend the Humean argument from attack at this point, we need to say what goes wrong with the example. I will offer a somewhat disjunctive response, with both disjuncts turning on the interest-relative account of justified belief that I offer in "Can We Do Without Pragmatic Encroachment?". I'll argue on the one hand that philosophers have been too quick to accept that we do not know we'll lose lotteries. As David Lewis pointed out, in many contexts it seems perfectly reasonable to say that people do have such knowledge. I'll argue that it often sounds right to say that because it's often true. On the other hand, I'll argue that in those settings where we do not know that the ticket will lose, c's testimony does not help us gain knowledge.

4 Interest-Relativity and Justification

In "Can We Do Without Pragmatic Encroachment?" I defended an interest-relative theory of belief. This implied an interest-relative theory of justified belief, even though the theory of justification was not, fundamentally, interest-relative. Rather, that theory held that what it was to justifiably believe that p was to have a high enough credence to believe p, and for that credence to be justified. What is 'high enough'? That, I claimed, was interest-relative. The agent's credence in p is high enough for her to believe p if her attitudes conditional on p matched her unconditional attitudes on every issue that was relevant to her. In particular, I said that for her to believe p, then for any A and B where the choice between doing A and B is a live question, and U is her utility function, then $[U(A) > U(B)] \leftrightarrow [U(A|p) > U(B|p)]$.

In that paper I also noted that sometimes the theoretical interests of the agent could be relevant to what she knows, but I don't think I went far enough down that road. Here's what I should have said. The idea behind my theory was that if you believe p, taking p as given in any inquiry doesn't change the results of that inquiry. If you believe p, you've already factored it in. Now one of the things that we can inquire into is the evidential probability of certain propositions given our evidence. If we already believe p, the results of those inquiries shouldn't change when we conditionalise on p. In particular, we should have the following two constraints on belief that p.

- If whether q is more probable than x is a live question, then $Cr(q) > x \leftrightarrow Cr(q|p) > x$.
- If the comparative probability of r and s is a live question, then $Cr(r) > Cr(s) \leftrightarrow Cr(r|p) > Cr(s|p)$.

The restriction to live questions here is important. If our credence in p is less than 1, even marginally less than 1, then there will be some inquiries whose results are altered by conditionalising on p. For instance, the question of whether p's probability is or isn't exactly 1 will be affected by whether we conditionalise on p. But that doesn't mean that belief requires probability 1. It means that not all inquiries are relevant to all agents, and in particular, the question of whether p's credence is exactly 1 isn't always relevant.

But consider one special case. Assume the agent *is* interested in just what the probability of p is. That is, for all x, the question of whether $Pr(p) > x$ is live for her. And assume that she judges that probability, on her evidence, to be less than 1. Assume also that she's rational enough to know that $Pr(p|p) = 1$. Then she can't believe that p, because there will be some x such that $Pr(p) < x$, but $Pr(p|p) > x$, and whether $Pr(p) > x$ is live.

I think that's a quite nice result. When we're trying to say what the relation is between credence and outright belief, it is tempting for many reasons to say that belief requires credence 1. One reason for that is that if we know the objective chance of p, and it's less than 1, it can feel very odd to say, without qualification, that we believe that p. It's much better to say that we believe p is probable. But it's very implausible to say that in general belief requires credence 1, because that would mean we believe very little.

The interest-relative view makes sense of this conundrum. On the one hand, belief does not in general require credence 1. On the other hand, when the agent is themselves focussed on the probability of p, they must judge that probability to be 1 to outright believe that p. I think that's a nice way to steer between the conflicting intuitions here.

5 Sources of Knowledge

From what I've said in the previous section, it's probably easy to imagine what I'll say about the challenge to (LDLC). The idea behind the challenge was two-fold. First, purely probabilistic evidence is not enough for knowledge. Second, other sources of evidence, such as testimony, can be the basis for knowledge even if we would, if pressed, say that they do not provide more support than purely probabilistic evidence. I'm going to accept the second claim (with some qualifications) but reject the first.

I think there are circumstances where we can, with Lewis, say the following.

> Pity poor Bill! He squanders all his spare cash on the pokies, the races, and the lottery. He will be a wage slave all his days. We know he will never be rich. (Lewis 1996: 443)

How, you might ask, can we know Bill will never be rich? The answer is that we know the odds are massively against him winning the lottery. That justifies a very high credence in his losing. For anything we care about, the odds are close enough to 1 that the difference doesn't matter. So our high credence is belief, and since it is justified, true, and undefeated, it is knowledge.

But wait, you say, isn't there some chance of Bill winning the lottery, and hence being rich? Why yes, there is. And doesn't that mean that we don't know he'll never be rich? Indeed it does. And doesn't that mean the previous paragraph is all mistaken? No, it doesn't. It means that asking all these questions changes the subject. In particular, it raises the question of whether the chance of Bill winning is equal to zero or greater than zero to salience. And once that question is salient, our degree of belief that Bill will lose is not close enough to 1 that the difference doesn't matter. The difference matters a lot, to the question you just raised. So I insist that given what I cared about a paragraph ago, I was speaking truly.

This explains why we think we cannot get knowledge on probabilistic grounds. Here's what we can't do. We can't simultaneously try to figure out what the probability of p is, conclude it is less than 1, and believe p. But that's simply because once the question of p's probability is live, we lose the belief that p. We can, I think, investigate whether the probability of p is, say, over 0.9, conclude that it is, and conclude on that basis that p.

As long as there are no further questions whose answer turns on whether p's probability is 1 or a little less, that could be enough for knowledge.

The converse is true about testimony. It's true that we can gain knowledge from testimony. And it's true that, if pressed, we may admit that that testimony is less than perfectly reliable. But what I deny we can do is admit the unreliability, work on figuring out just how unreliable it is, and hold onto the knowledge gained from testimony. But it's fairly intuitive that this would be impossible. Simultaneously thinking that my only reason for believing p is that S told me that p, and holding that S is somewhat unreliable, and may have been mistaken on this occasion, but nevertheless simply believing p, is an unstable state.

The difference between probabilistic grounds for belief, as when we believe we'll lose the lottery, and testimonial grounds then is not that one of them requires higher standards. It is rather that when we use explicitly probabilistic grounds, we tend to make probabilistic questions salient. And the salience of those questions destroys belief, and hence destroys knowledge. If we make the same questions salient in both the probabilistic and testimonial cases, we get the same criteria for knowledge. Hence the kind of case we've been considering is not a threat to (LDLC). Indeed, it is hard to see what could be a threat to (LDLC), without changing the salience of probabilistic questions. So I think (LDLC) survives, and anyone who wants to resist the Humean conclusion will have to look elsewhere to find the weak link in the argument.

6 Interest-Relativity and Knowledge

So far I've argued that given an interest-relative view of knowledge, we have a good defence of (LDLC) from putative counterexamples. But I haven't said what such an interest-relative view should look like, nor why we should endorse it. This section describes (in very sketchy outline) such a view, the next section is the defence.

It is a common view that knowledge has different defeaters to justified belief. This idea will be at the heart of expanding our view of justification into an account of knowledge. Knowledge will be, on this view, undefeated justified true belief. Unlike the account of justification, the account of defeaters will be interest-relative. (This does undermine the position I took in "Can we Do Without Pragmatic Encroachment?".)

Here is the kind of case that motivates thinking that knowledge has its own distinctive defeaters. S is offered a bet with the following payoff structure.

- If $p \wedge q$ is true, then the bet wins \$1
- If $\neg p \wedge \neg q$ is true, then the bet loses \$10
- If $p \leftrightarrow q$ is false, then the bet returns nothing.

It is given that p and q are probabilistically independent; we'll only consider credal functions where that obtains. We'll also assume that there is no other question about p in which the agent is interested. Finally, we'll assume that $Cr(p) = Cr(q) = 0.9$, and that the agent's credence in p is perfectly justified. Given the above account of justified belief, it follows that the agent has a justified belief in p. Her attitude to p is justified. And conditionalising on p does not change her attitude towards the bet - she prefers taking the bet both absolutely and conditional on p. If p were false she clearly wouldn't know that p, but let's assume that p is true. Given all that, does she know p?

There are two circumstances in which we might think that she does not. First, assume that her attitude towards q is completely unjustified. The rational credence in q, given her evidence, is 0.1. So the bet she's offered is a very bad bet, although it is a good bet conditional on p. So although conditionalising on p does not change any relevant attitudes, if she was being rational it would change attitudes on a relevant question. So I'm inclined to hold in this case that she doesn't know p, though she does have a justified belief in p.

The second case is trickier. Say that her credence in q is justified, but in fact q is false. If she learned $\neg q$ she clearly would think the bet is bad. But if she learned $\neg q$, then learned p, she wouldn't think the bet was bad; she would think it is irrelevant. So we might think that the fact that $\neg q$ is a defeater of her knowledge. I think that's at least sometimes correct. More formally, where F is a class of relevant facts, and A and B are salient options for the agent, for the agent to know that p, the following condition must be met.

- $(r \in F) \rightarrow [(U(A|r) > U(B|r)) \leftrightarrow (U(A|p \wedge r) > U(B|p \wedge r))]$

Intuitively, if r is relevant and true, then we should see whether conditionalising on p would make a difference once the agent has already conditionalised on r. If that's not the case, then r is a defeater for the agent's knowledge that p. As noted above, we should also look at theoretical interests as well as practical interests. So as well as that clause, we should have the following constraint on knowing that p whenever the question of whether s's probability is greater than x is relevant.

- $(r \in F) \rightarrow [(Pr(s|r) > x) \leftrightarrow [Pr(s|p \wedge r) > x]$

Again, the idea is that if r is relevant, then once the agent has conditionalised on r, conditionalising again on p shouldn't change the answer to a relevant question, in this case a question about the probability of s.

Just what is a relevant fact is a little hard to say, but some plausible claims about relevance suffice to show that this clause rules out certain cases of justified true belief that many people think are not knowledge. If we think that the agent's premises she has used in coming to be confident that p are relevant, and those premises are false, then these clauses won't be satisfied. So if the agent is confident in p only because she is very confident that $\neg r$, which supports p, then r is relevant. Since the question of whether $Pr(p) > 0.5$ is always relevant when the agent is thinking about p, and presumably $Pr(p|r) < 0.5$ if p's main support was $\neg r$, the above clause will be violated. That's how the interest-relative theory handles simple Gettier cases, for example.

There is a lot more to be said here. If we let any proposition be relevant, then any agent with any false belief will know very little, which would be wrong. What's needed is a systematic statement of relevance, and that's not something I have worked out at this stage. But hopefully I've said enough to suggest that an interest relative theory of belief, combined with an interest relative account of defeaters, is a coherent theory of knowledge. I'll conclude by arguing that the theory is more than coherent, it is the best solution to some hard problems in epistemology.

7 Motivating the Interest-Relative Theory

So far I've defended the a premise in the Humean sceptical argument by appeal to an interest-relative view about knowledge. And while I've spent a fair bit of time describing that interest-relative theory, I haven't spent that long deending it. The point of this section is to remedy that.

Unlike many others, I don't rely heavily on intuitions about cases to defend the interest-relative theory. My reason for declining to use this defence is not that the intuition-based arguments involving trains and banks and so on that are central to the literature fail. On the contrary, I find them quite convincing. But my experience is that others are not so moved by these cases, and so I want to try another route.

My preferred argument is a version of the argument for interest-relativity that runs through John Hawthorne's *Knowledge and Lotteries*. There seem to be several principles concerning knowledge and fallibility that any theory should account for, even though the principles are in some tension. Perhaps we won't be able to account for the truth of all of the principles, but it's a compulsory requirement on a theory of knowledge that it either show that the principles are approximately true, or explain why they seem so intuitive. Here are a list of the principles I think we should incorporate into our theory of knowledge.

Anti-Inductive Scepticism (AIS) We know a lot. Much of it is inductive, and hence goes beyond our evidence.

Conjunction Closure (CC) If we can know p, and we can know q, we don't need any more evidence to come to know $p \wedge q$.

Knowledge Provides Reasons (KPR) If we know p, and we know A is the thing to do given p, then our knowledge of p provides a good reason to do A.

Don't Bet on It (DBI) Sometimes, even though we know that p, it would be irrational to accept a bet that wins a small amount if p, but loses a vast amount if $\neg p$.

No Knowledge of Lotteries (NKL) We don't know what the result of a fair lottery will be, and in particular do not know that any given ticket will lose.

It's worth noting how platitudinous many of these points are. Hawthorne notes that the New York State Lottery uses as its slogan "You Never Know". They mean, to belabour the joke, that if you buy a ticket, you don't know whether it will win, i.e. (NKL). I suspect everyone will be on board with (AIS), but it's worth recalling Lewis's statement of similar sentiments at the start of "Elusive Knowledge".

> I know what food penguins eat. I know that phones used to ring, but nowadays squeal, when someone calls up. I know that Essendon won the 1993 Grand Final. I know that here is a hand, and here is another. We have all sorts of everyday knowledge, and we have it in abundance. To doubt that would be absurd.

I suspect (CC) is more controversial, especially since many people have argued that ditching (CC) is central to a solution to the preface paradox. But it is worth recalling just how absurd it would be to reject most instances of it. Imagine a philosopher who was interested in whether $p \wedge q$ was true. Our philosopher puts forward some good arguments for p. She then puts forward some good arguments for q. She then says, "Clearly we have more work to do to determine whether $p \wedge q$." That would be absurd! Nobody does that. Not even the critics of (CC). I've never seen a single theorist who thinks they have a further question to ask about whether the conjunction is true after working out the truth value of the conjuncts. There simply isn't a further question about whether $p \wedge q$ once we've settled whether p and settled whether q.

The real challenge concerns (KPR) and (DBI). I haven't yet formalised them carefully, and with good reason. Both principles seem extremely platitudinous. Something about each of them must be right. But the principles seem simply inconsistent. (KPR) says that we can use our knowledge in practical reasoning. So if we know p, and know that *A* is a good thing to do given *p*, we can do *A*. But (DBI) says that in some cases we can't do that; we can't bet on *p* even if we know it. One thing to say here is that common sense is simply inconsistent, and we have to abandon part of what it says. A better approach, I think, is to find a theory that rescues part of each platitude.

The interest-relative theory promises to do just that. The most impressive feature of the theory is how it reconciles (KPR) and (DBI). To explain this we need only assume, in addition to this theory, that knowledge entails justified belief. Whenever the agent does have to choose between *A* and *B*, and she knows *p*, she can use the fact that the conditional expected utility of *A* is greater than that of *B*, because given she knows *p*, that fact about conditional utilities will imply that the actual expected utility of *A* is greater than that of *B*. On the other hand, when she is thinking hypothetically about crazy bets on *p*, she can't assume that conditional and unconditional utilities match up, so she can't use her knowledge that *p*. If she were to be offered those bets, what is live would change, and hence so would what she knows. This seems like just the right result to me.

I won't say much about (CC) here, because the explanation of how this theory preserves (CC) is difficult, and I go over it at some length in the earlier paper. The explanation of (NKL) is quite simple, as we've seen above. Presumably when you are thinking about a lottery ticket, the question of what its probability of winning is becomes quite salient. That probability is greater than zero. If it were, we'd all agree you know it will lose. But given that you'll lose, the probability of winning is zero. So conditionalising on the claim that you'll lose does change the answer to some question you care about, namely the question of whether the probability of winning is zero, or greater than zero. So in any circumstance where you are thinking about a lottery ticket, you don't know it will lose.

The final issue is (AIS). This is a little harder, because fully evaluating what we know depends on getting clear about the role of defeaters, and that isn't something that is yet worked out. However we work things out, we have a lot of justified true beliefs on the interest-relative model. For most of our beliefs, there is no bet riding on them at such long odds that we need to worry about whether that belief is really well enough supported by the evidence. And the question of how many 9's there are at the start of the decimal expansion of the probability that that belief is true simply isn't a live question. So, assuming as I do that we are justified in having very high credences in a lot of propositions, we have a lot of justified beliefs.

The issue then comes down to the defeaters, and here it is harder to be so unequivocal. But note that in any area where (a) not too many of our beliefs are unjustified and (b) not too many of our beliefs are false, our justified beliefs will amount to knowledge. So I think the interest-relative theory implies that we know a lot, as (AIS) requires.