

Probability and Scepticism

Brian Weatherson

2014

One way to motivate scepticism is by looking at the ways we might possibly know we aren't brains in vats. Could we know we aren't brains in vats a priori? Many will say no, since it is possible to be a brain in a vat. Could we know it on the basis of evidence? The chapter argues that given some commonly held assumptions, the answer is no. In particular, there is a kind of sceptical hypothesis whose probability is decreased by conditionalising on the evidence we have. Using this fact, I argue that if we want to say our knowledge that we aren't brains in vats is a posteriori, we have to give up the view that all updating on evidence is by conditionalisation.

1 The Humean Sceptical Argument

The following, broadly Humean, sceptical argument is fascinating for many reasons.¹ In the argument E is the agent's evidence, H is some hypothesis derived by ampliative reasoning from her evidence, and \supset is the (classical) material conditional, i.e., $\neg E \vee H$.²

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 if the agent does know H , then it is possible for her to know $E \supset H$ by using \vee -introduction, 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 perhaps most strongly with Timothy Williamson's view that our evidence just is as our knowledge, is to accept the conclusion but deny it is as sceptical as it first appears (Williamson 1998, 2000). The Humean

¹On how closely this argument resembles Hume's argument for inductive scepticism, see Okasha (2001, 2005). I've previously discussed the argument in Weatherson (2005b) and Weatherson (2007).

²I'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.

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 even 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.

I'm going to argue here that such an epistemology has to deviate in one way or another from orthodox views. In particular, I'll argue that it has to accept deeply contingent *a priori* knowledge, or reject the idea that probabilistic updating should always go by conditionalisation.

2 A Probabilistic Argument for the *a posteriori* Premise

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.

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 false can only be known *a posteriori*.³ 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.⁴

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.

Credences are Classical Probabilities (CCP) Cr is a classical probability function.

Updating Theorem (UT) Let $E = E_1 \wedge \dots \wedge E_n$, $\Pr(E) > 0$, $\Pr(H) < 1$, $\Pr(E \supset H) < 1$ and for each i , $\Pr(E_i) < 1$. And assume \Pr is a classical probability function. Then $\Pr(E \supset H \mid E_i) < \Pr(E \supset H)$.

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 $Cr_Y(X) = Cr(X \mid Y)$ for all X, Y .

³I'm using 'could turn out to be false' in the sense described in Yablo (2002).

⁴When I say it is an 'empiricist' assumption, I mean that the two ways of rejecting it correspond to two things classical empiricists rejected. One is that we can reason our way, perhaps abductively, to substantive knowledge about the external world, a la Vogel (1990) or Bonjour (1997). The other is that we have substantial innate knowledge about the external world, and this is not justified by empirical evidence, but perhaps by its reliability. It's interesting that these two forms of rejection are associated with very different views in contemporary philosophy, but they seem both anti-empiricist to me.

Learning Doesn't Lower Credence (LDLC) It is impossible for a rational agent to learn X on the basis of evidence Y if $Cr_Y(X) < Cr(X)$.

Knowledge Requires Learning (KRL) If the agent knows $E \supset H$ *a posteriori*, i.e., on the basis of her empirical evidence, then there is some part of her evidence E_i on the basis of which she learned $E \supset H$, and before she learned it, her credence in E_i is less than 1.⁵

Humean Conclusion (HC) So, it is impossible to know $E \supset H$ *a posteriori*.⁶

Now if someone wants to reject the Humean argument at premise 2, they better reject one of these five principles. But the principles are each reasonably strong.

2.1 Classical and Non-Classical Credences

There is a huge literature on whether credence functions should be probability functions. For a good recent overview, see Hájek (2008). Most of that literature has assumed that the underlying logic we use in reasoning under uncertainty should be classical. But this assumption can be questioned too, as I did in Weatherson (2003). It turns out this matters for the argument here. Without some extensive use of classical assumptions, it doesn't always hold that $\Pr(E \supset H|E) < \Pr(E \supset H)$. (For more on this, see Jehle and Weatherson (2012).) In principle, that's one possible way out of the argument. But I imagine it will be too costly a way out for most philosophers.

2.2 The Updating Theorem

This is a theorem, so it is harder to reject! It's not a new theorem by any stretch; in fact it is a fairly simple result. Here's a proof of it. The proof uses the following very familiar result from the classical probability calculus.

$$\Pr(X) = \Pr(X | Y)\Pr(Y) + \Pr(X | \neg Y)\Pr(\neg Y)$$

We'll substitute $E \supset H$ for X and E_i for Y , to get.

$$\Pr(E \supset H) = \Pr(E \supset H | E_i)\Pr(E_i) + \Pr(E \supset H | \neg E_i)\Pr(\neg E_i)$$

⁵When I say 'part' here, I just mean that E_i is one of the conjuncts of E . This may require relabelling, if for instance the basis for $E \supset H$ consists of many conjuncts of E under some representation; just collect all those into a single conjunct.

⁶*Proof:* Assume for *reductio* that we can know $E \supset H$ *a posteriori*. So by (KRL) there is some E_i that is the basis of this knowledge, such that before she learned it, her credence in it was less than 1. When she does learn E_i , she conditionalises her credences, as required by (UIC). So updating (i.e., conditionalising) on E_i raised its credence to 1, but by (LDLC) did not lower the agent's credence in $E \supset H$. As we noted in footnote 2, we're assuming $Cr(E) > 0$, and by (CCP) we're assuming (Cr) is a classical probability function, so $Cr(E_i) > 0$. We also assumed $Cr(H) < 1$. So the conditions for applying (UIT) are all satisfied, and hence her credence in $E \supset H$ goes down when she updates in E_i . That contradicts our earlier conclusion that it does not go down, completing the *reductio*.

Since $\neg E_i$ entails $E \supset H$ it follows that $\Pr(E \supset H | \neg E_i) = 1$. And since $\Pr(E \supset H | E_i) \leq 1$, it follows that $\Pr(E \supset H | E_i) \leq \Pr(E \supset H | \neg E_i)$. So substituting $\Pr(E \supset H | E_i)$ for $\Pr(E \supset H | \neg E_i)$ on the right-hand side of that equation, and noting that we can't make the right-hand side larger by that substitution, we get,

$$\Pr(E \supset H) \geq \Pr(E \supset H | E_i)\Pr(E_i) + \Pr(E \supset H | E_i)\Pr(\neg E_i)$$

with equality only if $\Pr(E \supset H | E_i) = \Pr(E \supset H | \neg E_i) = 1$ or $\Pr(\neg E_i) = 0$. But we assumed that $\Pr(E_i) < 1$, so $\Pr(\neg E_i) \neq 0$. We'll come back to the argument that $\Pr(E \supset H | E_i) < 1$. Note for now that we can rewrite that inequality by factoring out $\Pr(E \supset H | E_i)$, to get

$$\Pr(E \supset H) \geq \Pr(E \supset H | E_i)(\Pr(E_i) + \Pr(\neg E_i))$$

But $\Pr(E_i) + \Pr(\neg E_i) = 1$, is a trivial theorem of the classical probability calculus, so this just reduces to:

$$\Pr(E \supset H) \geq \Pr(E \supset H | E_i)$$

Since $\Pr(E \supset H) \geq \Pr(E \supset H | E_i)$, and we assumed $\Pr(E \supset H) < 1$, it follows that $\Pr(E \supset H | E_i) < 1$. But that means neither condition for the inequality introduced above to not be strict is satisfied. So in fact we can conclude:

$$\Pr(E \supset H) > \Pr(E \supset H | E_i)$$

as required.

2.3 Updating and Conditionalising

In Weatherson (2007), I argue that philosophers who are sympathetic to empiricism (broadly construed) should reject (UIC). That's because (UIC) embodies a very implausible picture of the relationship between evidence and hypotheses. We can see this more clearly if we think about the non-probabilistic case first. Consider the following hypothesis.

- After learning E , an agent should believe H iff they believed $E \supset H$ before learning E .

This picture suggests that all a rational agent has to do is line up all their thoughts at the beginning of time, or I guess of inquiry, and then go around collecting evidence and applying modus ponens. Indeed, it says there is nothing else that would be rational to do. This strikes me as implausible in the extreme. There are many more rules we can use to get from evidence to conclusion than modus ponens applied to pre-known conditionals. Sometimes, it is only by getting some evidence that we are in a position to see what that evidence supports.⁷

⁷In (Weatherson 2007) I argue for this by considering agents with radically different kinds of evidence to ours, and noting how much we could know about what kinds of conclusions their evidence supports, and what they could know about what kinds of conclusions our evidence supports.

Now the rule that we should always update by conditionalisation is like the rule that we should always update by modus ponens in the way just suggested. Instead of saying that learning E doesn't change which conditionals with antecedent E we can know to be true, it says that learning E doesn't change the conditional probability of anything given E . And it seems equally implausible for just the same reason.

So I don't think (UIC) is right, and I suspect at the end of the day rejecting it is the best way to avoid the Humean sceptical argument. But I do think that there are many people who are not so sceptical (as a casual perusal of the literature on conditionalisation will show). And there may be several others who are implicitly committed to (UIC), whether or not they explicitly acknowledge that fact. So I think it is interesting to see how (UIC) can promote a certain kind of scepticism.⁸

2.4 Learning and Credence

We will look at a lot of cases that seem to raise problems for (LDLC) below. But first I just wanted to note that the restriction to rational agents avoids one quick problem for the principle. An irrational agent might simply ignore very good evidence for p , and then come to believe p on the basis of evidence that undermines that initial evidence for p , but provides an independent (but weaker) basis for believing p . She really could learn p on the second basis, even though its probability was lowered.

The restriction to rational agents is intended to rule out such a case. We assume that the agent has antecedently taken correct account of the available evidence. If that isn't the case, then something which lowers the probability of p can ground knowledge that p , perhaps because it reinforces evidence that S had, but was not properly using. What's interesting is whether we can have violations of (LDLC) without irrationality.

2.5 Learning and Knowing

You might think that the last premise, (KRL), would be the easiest one to defend. Arguably something even stronger is an analytic truth, namely that S learns p at t iff S knows p at t but not before p . Indeed, I used to think this. But it isn't actually true. What is plausibly true, as we'll see by some reflections on learning, is that knowing requires either innate knowledge or learning. But the relationship between the learning and the knowing may be very complicated indeed. Let's turn to that relationship now.

3 Learning and Defeaters

In an earlier version of this paper, I worked with a much simpler premise, namely that coming to know required probability non-decrease. But that isn't right.⁹ The problem is that the view in question doesn't account for defeaters.

⁸Of course, I'm hardly the only person to promote doubts about (UIC). See Arntzenius (2003) for some very different kind of criticisms.

⁹The essential reason it isn't right was pointed out by Martin Smith in comments on this paper at the 2009 Arché scepticism conference. This section is basically a response to the good objections he raised to the earlier version of the paper.

Here's a schematic version of the kind of case that causes problems. Assume S has a justified true belief that p . Assume also that there is some defeater D that blocks S 's belief from being knowledge. Now imagine an event happens that (a) slightly lowers the evidential probability of p for S , and (b) defeats the defeater D . Then after the event, it may be that S knows that p , although she does so in part in virtue of an event that lowered the probability of p .

The schematic version of this argument is much more plausible than any particular case, since defeaters are often very hard to get clear judgments about. But here are three cases that may illustrate the kind of thing I have in mind.

Dead Dictator

Carol is trapped in Gilbert Harman's dead dictator story (Harman 1973, 75). At t_1 she reads the one newspaper that correctly (and sensitively) reported that the dictator has died. She hasn't seen the copious other reports that the dictator is alive, but the existence of those reports defeats her putative knowledge that the dictator is alive. At t_2 , all the other news sources change their tune, and acknowledge the dictator has died. Carol doesn't see any of those newspapers; she's too busy playing Farmville. But Carol's memory very slowly degrades over time (as most memories do), so at t_2 her evidence that the dictator died is slightly weaker than at t_1 . Still, over the time between t_1 and t_2 while she played Farmville, she came to know the dictator had died, even while the (evidential) probability of that decreased.

Fake Barns

Bob starts our story in Fake Barn Country (Goldman 1976). At t_1 , he starts looking straight at a genuine barn on a distant hill, and forms the belief that there is a barn on that hill. Since he's in fake barn country, he doesn't know there is a barn on the hill. At t_2 , while Bob is still looking at the one genuine barn, all the fake barns are instantly destroyed by a visiting spaceship, from a race which doesn't put up with nonsense like fake barns. The mist from the vaporised barns slightly clouds Bob's vision, so he doesn't have quite as clear a view of the barn on the hill. But he still has an excellent view, so after the barns are destroyed, Bob's belief that there is a barn on that hill is knowledge. So at t_2 he comes to know, for the first time, that there is a barn on that hill. But the vaporisation of the fake barns, which is what lets him come to know that there is a barn on that hill, doesn't raise the (evidential) probability that there is a barn there.¹⁰ Indeed, by making Bob's vision a little cloudier, it lowers that probability.

Gettier Cases

Ted starts our story believing (truly, at least in the world of the story) that Bertrand Russell was the last analytic philosopher to win the Nobel

¹⁰It does raise the probability that a randomly selected barn-like structure in Bob's vicinity is a barn, but that's not the evidential probability for Bob of there being a barn in that hill.

Prize in literature. The next day, the 2011 Nobel Prize in literature is announced. At t_1 , a trustworthy and very reliable friend of Ted's tells him that Fred has won the Nobel Prize in literature. Ted believes this, and since Fred is an analytic philosopher, Ted reasonably infers that, as of 2011 at least, Bertrand Russell was not the last analytic philosopher to win the Nobel Prize in literature. This conclusion is true, but not because Fred won. In fact, Ed, who is also an analytic philosopher, won the 2011 Nobel Prize in literature. At t_2 , Ted is told by a friend who is just slightly less reliable than the first friend that it is Ed, not Fred, who won the prize.¹¹ Since Ted knows that Ed is also an analytic philosopher, this doesn't change his belief that Bertrand Russell was not the last analytic philosopher to win the Nobel Prize in literature. But it does change that belief from a mere justified true belief into knowledge.

At t_1 , Ted didn't know that Bertrand Russell was not the last analytic philosopher to win the Nobel Prize in literature, since his true belief was based on a falsehood.¹² At t_2 , he did know this, on the basis of the second friend's testimony. But since the second friend was less reliable, and since the second piece of testimony raised doubts about the first in ways that render each of them suspect, the probability of Ted's conclusion was lower at t_2 than t_1 . So the second piece of testimony both lowered the probability of Ted's conclusion, and turned it into knowledge.

In every one of those cases, something happens that ever so slightly lowers the probability of p , and also defeats a defeater of the agent's knowledge that p . So the agent gets knowledge that p in virtue of an event that lowers the probability of p .

But there is, in general, something odd about events that bring about a conclusion by double prevention. There's a big difference between being responsible for a pot of soup in virtue of preparing and cooking it, and being responsible for it in virtue of removing the banana peel that the chef would have slipped on when bringing the pot to the table. The same goes for knowledge; things that remove defeaters of knowledge are importantly different in kind to the underlying bases for knowledge.

The difference in question is one that we mark in language. We say that the chef cooked, or prepared, the soup. We don't say that the banana peel remover did either of those things, although she may have caused the soup to be ready to eat. In the three cases described above, I think it's odd to say that the agent *learns* that p in virtue of the defeater being defeated.¹³

Carol can't learn that the dictator has died while she is busy playing Farmville, and not being in any contact (of the right kind) with the outside world. So the passage of

¹¹ Presumably for (Gettier 1963).

¹² I'm not presupposing here that we can never get knowledge from false beliefs, just that the falsity of Ted's initial belief explains why his subsequent belief is not knowledge. For more on this point, see Warfield (2005).

¹³ A quick sample of informants suggests that this is much less odd in the Gettier case than in the other two cases. We'll come back to this point below.

time from t_1 to t_2 doesn't cause her to *learn* the dictator has died. If she ever learned this, she learned it at t_1 . And surely she did learn it. It wasn't innate knowledge, and it wasn't knowledge that was somehow implanted in her, in the way characters in the movie *The Matrix* can have knowledge implanted directly into their brain.¹⁴ So she learned the dictator died, and the only learning she did took place at t_1 , so she learned that the dictator died at t_1 .

I think the same thing is true in the other cases. Bob learns that there is a barn on that hill at t_1 , but doesn't know this until t_2 . And Ted learns that Russell is not the last analytic philosopher to win at t_1 , but doesn't know this until t_2 . So actually cases where defeaters are defeated by probability lowerers are not counterexamples to (LDLC).

Officially, that completes my defence of (LDLC) from this kind of objection. But I know that not everyone agrees with my judgments about these three cases, especially the last. So I wanted to say a bit about why the overall argument is not overly affected even if I'm wrong about (LDLC).

Note that in all three of the cases, there are two distinctive things that happen at t_1 . The agent gets a lot of evidence in favour of p . And the agent gets some kind of defeater that prevents beliefs based on that evidence turning into knowledge. Now let's say that the probabilistic argument that $E \supset H$ can't be known *a posteriori* fails because of an analogy with these cases. That is, let's suppose that $E \supset H$ can be known *a posteriori* even though all the empirical evidence *lowers* its probability, and the explanation for how this is possible is by analogy with cases like **Dead Dictator**. Then we should be able to find analogies for these two properties: something sometime raises the probability of p , and there is a defeater that prevents p being known despite having a high probability.

The first putative point of analogy obviously fails. After all, $E \supset H$ was designed so that the agent *never* gets evidence that raises its probability. So we should already be suspicious of such an analogy going through. But the second putative point of analogy is actually pretty interesting. Could there be a defeater that prevents someone *knowing a priori* that $E \supset H$ even though the *a priori* probability of $E \supset H$ is very high?¹⁵

I don't have a conclusive argument that there is no such defeater, but it's worth noting that most of the usual suspects don't seem to work.

Sensitivity It's true that the *a priori* belief that $E \supset H$ is insensitive. That is, even if it were false, it would still be held. But the *a posteriori* belief that $E \supset H$ is also insensitive. So if insensitivity is a barrier to knowledge, this is a quick argument for the conclusion of the Humean sceptical argument, not a way to block a premise in an argument for premise 2.¹⁶

¹⁴It's a delicate question whether this kind of procedure is properly called learning. I'm inclined to say that it is, but I suspect a lot of people aren't, so didn't want to presuppose my own idiosyncratic usage here. Thanks here to Jonathan Livengood and Daniele Sgaravatti.

¹⁵Why are we interested in whether we can prevent *a priori* knowledge of $E \supset H$? Because we're interested in ways in which $E \supset H$ can be known *a posteriori*, and by definition that means that it isn't known *a priori*. The idea I'm floating here, which I don't think will work, is that the first *knowledge* of $E \supset H$ is after the agent gets some evidence, and because she gets that evidence, although $E \supset H$ has maximal probability *a priori*, i.e., before she gets any evidence.

¹⁶Vogel (1987) makes a similar point that sensitivity and induction don't mix.

Safety The belief that $E \supset H$ is true seems to be safe. After all, any world in which it is false must be rather distant. If not, then we don't know very much about the external world, which means we have a direct argument for the conclusion of the Humean sceptical conclusion, not a way to block a premise in an argument for premise 2.

Reliability There are a few reliable ways in which $E \supset H$ could be believed. One is the rule, in any circumstance, believe $E \supset H$. More practically, the rule that says whenever X is good evidence for Y , good enough to ground knowledge that Y , and one doesn't have any evidence for $X \wedge \neg Y$, then believe $X \supset Y$ seems fairly reliable too. So there isn't an obvious reliability argument that $E \supset H$ is not knowable *a priori*.

False Belief It's possible to infer $E \supset H$ *a priori* from a false premise. But it isn't necessary. The inference from the premise that E is good evidence for H to the conclusion $E \supset H$ seems reasonable, and based on true (indeed knowable) premises.

In short, the following position looks untenable to me: It's possible to have *a priori* a justified true belief in $E \supset H$, but defeaters always conspire to ensure that this cannot rise to the level of knowledge. There just aren't the defeaters around to ensure this works.

A corollary to this is that it is impossible to learn $E \supset H$ on the basis of a probability lowerer that simultaneously defeats an *a priori* defeater to $E \supset H$. There just aren't enough defeaters around for that strategy to work.

4 Learning, Probability and Interests

A slightly different kind of objection to (LDLC) comes from considerations about lottery cases. My reply, in short, is going to be that standard treatments of lottery cases are not very promising, that we should adopt a kind of interest-relative approach to lottery cases instead, and when we do that the problem goes away. But first I'll set out the problem.¹⁷

4.1 Lotteries and Learning

The case we will focus on 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 .

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

¹⁷I'm grateful to David Chalmers, Crispin Wright and Elia Zardini for pressing me on the need to address these cases. The cases are discussed in more detail in Smith (2010) and Zardini (2012).

That is, testimony from a knowledgeable source is maximally valuable testimony, while testimony from other sources has no evidential value. The second assumption is a little extreme¹⁸, but more moderate models will also generate the kind of example we're interested in here.¹⁹

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. The agent wins iff 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 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 is certain that what b says is true. And her credence in Rc is 0.99, which is reasonably high. (But we'll come back to the question of just how high it is by everyday standards.) Still, 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

¹⁸I'm interpreting R in such a way that Ra entails what a says is true, so $Ra \wedge Sap$ entails p , so the first assumption is natural. Making the second assumption more realistic would just increase the complexity of the model without revealing anything insightful. Since this model is meant to raise problems for my view, I think it is fine to use an extreme case, and not complain about its extremity.

¹⁹I 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.

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 when our credence in the reliability of the testifier 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 by increasing the number of coin flips, we can make the credence even higher; arbitrarily close to 1 if need be.) 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.

The first thing to note about this example is 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.

It isn't at all clear that anything similar is happening in the case of E grounding knowledge of $E \supset H$. After all, the point of the theorem we earlier proved was that *every single part* of E lowers the probability of $E \supset H$. Now I don't want to rest too much on a theory of how evidence divides into parts, and maybe there won't be any way to make sense of the notion of parts of evidence in a way that is needed for the point I'm making here to work. If we are to have a theory of parts of evidence, I like a causal theory of evidence that naturally lends itself to individuating parts as being evidence that arrives via different causal chains. But I don't think we know nearly enough about the ontology of evidence to make this kind of response compelling.

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 defend in Weatherson (2005a) and Weatherson (2012). 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 (1996) 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.

4.2 Interest-Relativity, Knowledge and Justification

In Weatherson (2005a) 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 (in a sense I describe in much more detail in the earlier paper), 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.²⁰ 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. 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 exactly 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

²⁰I go into much more detail on this in Weatherson (2012).

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.

Given all this, 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.²¹

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.²²

I think that what's going on in cases like these involves the interest-relativity of belief, not in the first-instance the interest-relativity of knowledge. Does that mean that if an agent held on to their beliefs across changes of interest, then their knowledge would not be affected by changes of interest? No; because the only way to hold on to beliefs when interests change may involve raising one's credence so high that it would be irrational, and when credences are irrational the resulting beliefs are irrational, and irrational beliefs can't constitute knowledge.

My positive view is a form of interest-relative *invariantism*; that is, I don't think contextualism is true about 'knows'. But I haven't relied on that here, just on the interest-relativity. If one wanted to hold a form of interest-relative *contextualism*, a la Fantl and

²¹ So I'm disagreeing with those such as Nelkin (2000) who think high probability can't suffice for knowledge.

But I think the comments below help explain away the motivations for such views.

²² Two technical points about how what I said relates to the broader debates about interest-relativity.

McGrath (2009), this explanation would still go through. There are puzzles that might push one towards interest-relative *contextualism*, but I think there are larger puzzles that should push one back towards invariantism (Weatherson 2006).

This explains why we think we can't 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 hence live.²³ 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.

Here's a crude summary of these reflections. If questions about the precise probability of H or $E \supset H$ are salient, then $E \supset H$ can't be known before or after learning E . If they aren't, $E \supset H$ can be known both *a priori* and *a posteriori*. The only way we get that $E \supset H$ is only knowable after learning E is if we equivocate between the two positions on what is salient.

5 Conclusions

So I think (LDLC) is invulnerable to these kinds of objections. Since it is intuitively a very plausible principle, and these attempts to counterexample it have failed, I think we should adopt as a working hypothesis that it is true. That means, I think, that we really have two options for responding to the Humean argument.

²³Salient to the person doing the reasoning that is. As an invariantist, I think that matters. But a contextualist who thought what's relevant to subjects is thereby relevant could say the same thing.

1. Accept that $E \supset H$ is *a priori* knowable.
2. Reject (UIC), and say some updating is not by conditionalisation.

I don't think either of these are bad options. You can read Weatherson (2005b) as an attempt to defend the first, and Weatherson (2007) as an attempt to defend the second. But I do think these options aren't available to everyone.

If $E \supset H$ is *a priori* knowable, then any kind of 'modal' account of the *a priori* has to fail. That is, we can't understand *a priori* as any kind of metaphysical necessity, since $E \wedge \neg H$ is clearly possible.²⁴ It's just that we have defeasible, fallible *a priori* knowledge that it isn't true. And I noted above that (UIC) will follow from some other independently attractive views about what we can know *a priori* about epistemology, and when it is that conditionalising seems wrong. Many years ago, I held both (UIC) and that deeply contingent truths like $E \supset H$ could not be known *a priori*. I now think that's an unstable combination of views; it leaves you without resources to turn back the Humean argument.

References

- Arntzenius, Frank. 2003. "Some Problems for Conditionalization and Reflection." *Journal of Philosophy* 100 (7): 356–70. doi: 10.5840/jphil2003100729.
- BonJour, Laurence. 1997. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Fantl, Jeremy, and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–23. doi: 10.2307/3326922.
- Goldman, Alvin I. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73 (20): 771–91. doi: 10.2307/2025679.
- Hájek, Alan. 2008. "Arguments for-or Against-Probabilism?" *British Journal for the Philosophy of Science* 59 (4): 793–819. doi: 10.1093/bjps/axn045.
- Harman, Gilbert. 1973. *Thought*. Princeton: Princeton University Press.
- Jehle, David, and Brian Weatherson. 2012. "Dogmatism, Probability and Logical Uncertainty." In *New Waves in Philosophical Logic*, edited by Greg Restall and Gillian Russell, 95–111. Basingstoke: Palgrave Macmillan.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549–67. doi: 10.1080/00048409612347521. Reprinted in his *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999, 418–446. References to reprint.
- Nelkin, Dana. 2000. "The Lottery Paradox, Knowledge, and Rationality." *Philosophical Review* 109 (3): 373–409. doi: 10.2307/2693695.
- Okasha, Samir. 2001. "What Did Hume Really Show about Induction?" *The Philosophical Quarterly* 51 (204): 307–27. doi: 10.1111/1467-9213.00231.

²⁴I mean both that it's true in some possible worlds, and in some worlds considered as actual, so a 'two-dimensional' equation of a priority with a kind of metaphysical possible is ruled out.

- . 2005. “Does Hume’s Argument Against Induction Rest on a Quantifier-Shift Fallacy?” *Proceedings of the Aristotelian Society* 105 (1): 237–55. doi: 10.1111/j.0066-7373.2004.00113.x.
- Smith, Martin. 2010. “What Else Justification Could Be.” *Noûs* 44 (1): 10–31. doi: 10.1111/j.1468-0068.2009.00729.x.
- Vogel, Jonathan. 1987. “Tracking, Closure and Inductive Knowledge.” In *The Possibility of Knowledge: Nozick and His Critics*, edited by Stephen Luper-Foy, 197–215. Totowa, NJ: Rowman & Littlefield.
- . 1990. “Cartesian Skepticism and Inference to the Best Explanation.” *Journal of Philosophy* 87 (11): 658–66. doi: 10.5840/jphil1990871123.
- Warfield, Ted A. 2005. “Knowledge from Falsehood.” *Philosophical Perspectives* 19: 405–16. doi: 10.1111/j.1520-8583.2005.00067.x.
- Weatherson, Brian. 2003. “From Classical to Intuitionistic Probability.” *Notre Dame Journal of Formal Logic* 44 (2): 111–23. doi: 10.1305/ndjfl/1082637807.
- . 2005a. “Can We Do Without Pragmatic Encroachment?” *Philosophical Perspectives* 19 (1): 417–43. doi: 10.1111/j.1520-8583.2005.00068.x.
- . 2005b. “Scepticism, Rationalism and Externalism.” *Oxford Studies in Epistemology* 1: 311–31.
- . 2006. “Questioning Contextualism.” In *Epistemology Futures*, edited by Stephen Cade Hetherington, 133–47. Oxford: Oxford University Press.
- . 2007. “The Bayesian and the Dogmatist.” *Proceedings of the Aristotelian Society* 107: 169–85. doi: 10.1111/j.1467-9264.2007.00217.x.
- . 2012. “Knowledge, Bets and Interests.” In *Knowledge Ascriptions*, edited by Jessica Brown and Mikkel Gerken, 75–103. Oxford: Oxford University Press.
- Williamson, Timothy. 1998. “Conditionalizing on Knowledge.” *British Journal for the Philosophy of Science* 49 (1): 89–121. doi: 10.1093/bjps/49.1.89.
- . 2000. “Scepticism and Evidence.” *Philosophy and Phenomenological Research* 60 (3): 613–28. doi: 10.2307/2653819.
- Yablo, Stephen. 2002. “Coulda, Woulda, Shoulda.” In *Conceivability and Possibility*, edited by Tamar Szabó Gendler and John Hawthorne, 441–92. Oxford: Oxford University Press.
- Zardini, Elia. 2012. “Confirming the Less Likely, Discovering the Unknown Dogmatism: Surd and Doubly Surd, Natural, Flat and Doubly Flat.” In *Scepticism and Perceptual Justification*, edited by Dylan Dodd and Elia Zardini, 33–70. Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780199658343.001.0001.