

### **The Paradox of Infinite Distrust**

Roger White has drawn my attention to an interesting problem, having to do with what to believe in a situation in which you have evidence that the world is infinite. I will build up to the situation in stages.

#### Stage One

You sit before a table. On the table are a phone and a cigar box. Into the cigar box you throw a fair die. You close the lid of the box without looking at how your die landed.

How confident should you be that your die landed six?

#### Stage Two

Reliable authorities on things cosmological tell you that space and matter are infinite. They tell you that the universe contains infinitely many people in a position like yours – sitting before a table, on which there are a red telephone and a cigar box, into which you have just thrown a fair die. They tell you that infinitely many of these people got sixes (they call these people the *sixers*) and infinitely many of these people got non-sixes (they call these people the *non-sixers*). And they tell you that, as was to be expected, the non-sixers are distributed more densely across the universe than the sixers. For any point in space, the proportion of sixers to non-sixers in a sphere centered on that point tends towards one sixth as the diameter of the sphere tends towards infinity.

Now how confident should you be that your die landed six?

#### Stage Three

The authorities tell you that there is an interesting spatial point, the *central point of the universe*. Neither you nor any of the dice-throwers know where it is. They use this point to define a pairing relation: Let the *buddy* of the closest sixer to the central point be the

closest non-sixer to the central point. Let the *buddy* of the second closest sixer be the second closest non-sixer... and so on.

Now how confident should you be that your die landed six? How confident should you be that your buddy's die landed six?

#### Stage Four

The authorities (with awesome cellular networks at their disposal) tell you that they are enabling each person to talk with his or her buddy. The phone at your table rings. You pick it up. There, on Facetime, is your buddy's table. There is your buddy's cigar box. There is your buddy. You and your buddy talk and talk. Some bits of the conversation are surprising to you: Your buddy's name is Isa. She lives on a planet with two moons. But some bits of the conversation are exactly as you expected them to be: With respect to tables, authorities, cigar boxes, phones and dice, your buddy's situation is just like yours.

Now how confident should you be that your die landed six? How confident should you be that your buddy's die landed six?

Please think about these questions carefully. To make them seem pressing, I suggest that you imagine, at each stage, that something really important is at stake: the life of your dog. Imagine, at each stage, being forced to choose between these two deals.

*Deadly Six*                      Your dog dies iff your die landed six.

*Deadly One-Two*              Your dog dies iff your die landed one or two.

Which deal *would* you choose? Which deal *should* you choose?

## **2. A Paradox**

The following three claims are collectively inconsistent but *truthy* – each has the

appearance of truth.

*Start at 1/6*

At Stage One you should be 1/6 confident that your die landed six.

Why? Because it is a fair coin! There is nothing funny going on at Stage One.

*No Changes*

At every stage, if you should be 1/6 confident at that stage, then you should be 1/6 confident at the next stage.

Why? Intuitively: Because, as you transition to stages Two, Three and Four, you don't learn anything that bears on whether your die landed six. The infinity of the cosmos does not tell you anything about your die. The existence of the central point does not tell you anything about your die. Your conversation with your buddy does not tell you anything about your die. Absent learning something new about your die, your epistemic obligations stay the same.

*Symmetry*

At Stage Four you should be 1/2 confident that your die landed six.

Why? Intuitively: Because you and Isa are in just the same evidential situation! Everything she knows about her die, her cigar box, the background situation, you know. Everything you know about your die, your cigar box, the background situation, she knows. Neither of you have any secrets from one another on topics related to dice. Given that you have the same evidence, know all the same things, your levels of confidence should match. And, given the symmetry of the situation, the level at which they should match must be 1/2.

### 3. The Underlying Tension

Before we try to solve this paradox, let's go deeper, and in more detail, into the motivation for each of the conflicting claims. *Start at 1/6* is non-negotiable. To deny it is to give up on normative epistemology altogether. Yet each of the other two follows naturally from an established doctrine in normative epistemology.

*No Changes* follows naturally from the doctrine that proper transitions in confidence conform to *Bayesian Conditionalization*. Roughly, the idea is that when you are taking an interest in whether a proposition,  $p$ , is true, and you learn something else,  $q$ , your subsequent confidence in  $p$  should match your prior confidence in  $p$  on the supposition that  $q$ . In this case, at every stage your confidence that your die landed six on the supposition that things are as you learn them to be in the next stage should be  $1/6$ , so at every stage your confidence that your die landed six should remain at  $1/6$ .

It takes a bit of work to make this rough idea precise and apply it to our case. First we need to suppose that for some propositions  $p$ ,  $q$ , in addition to a level of confidence in  $p$  and a level of confidence in  $q$ , you have a level of confidence in  $p$  on the supposition that  $q$ . We call this your *conditional confidence* in  $p$  given  $q$ , or  $cr(p | q)$ .

Next we suppose that your conditional confidences bear a certain relation to your conditional confidences. For any propositions,  $p$ ,  $q$ :

*Conditional Confidence*<sup>1</sup>

$$cr(p | q) = cr(p \& q) / cr(q)$$

---

<sup>1</sup> I should mention that some Bayesians take this to be a *definition* of conditional confidence – to have the conditional attitude just is to have the pattern of unconditional attitudes. Others take it to express a necessary connection between distinct attitudes – necessarily, you have the conditional attitude iff you have the pattern of unconditional attitudes. Others take it merely to express a rational necessity – necessarily, if you are rational, then you have the conditional attitude iff you have the pattern of unconditional attitudes. For present purposes it won't matter which way we go here.

From this, plus the assumption that your credences conform to the probability calculus, we can derive:

*The No-Difference Principle*

If  $cr(p|q) = cr(p|\neg q)$ , then  $cr(p|q) = cr(p)$

In words: If, so far as your conditional confidence in  $p$  goes, it makes no difference whether you suppose that  $q$  is true or false, then your conditional confidence that  $p$  on the supposition that  $q$  is your unconditional confidence that  $p$ .

Next we suppose an epistemic norm. Here's a first pass at it. For any propositions  $p$ ,  $q$ , level of confidence  $x$

*Conditional Updating (1)*

If, prior to learning that  $q$ , your confidence in  $p$  on the supposition that  $q$  is  $x$ , then, upon learning that  $p$ , your confidence in  $p$  should be  $x$ .

But this won't quite do. In some ways it is too strong: What if it is not the case that, prior to learning that  $q$ , your confidence in  $p|q$  should be  $x$  (either because you are making a mistake in having confidence  $x$  in  $p|q$ , or because it is a situation in which many different degrees of confidence would be okay)? Surely then it is not the case that, after learning that  $q$ , your confidence in  $p$  should be  $x$ . In other ways it is too weak properly to capture the Bayesian idea. What if, prior to learning that  $q$ , your confidence in  $p|q$  should be  $x$ , but is not  $x$ ? Surely then, after learning that  $q$ , your confidence in  $p$  should be  $x$ . Here's a better principle:

*Conditional Updating (2)*

If, prior to learning that  $q$ , your confidence in  $p$  on the supposition that  $q$  should be  $x$ , then, upon learning that  $q$ , your confidence in  $p$  should be  $x$ .

Though better, it is still too strong. What if there are further relevant changes to your doxastic state, as you learn that  $p$ ? What if  $q$  is not the strongest thing you learn? What if you suffer some form of *doxastic corrosion* – what if, for example, you forget things<sup>2</sup>, or gain reasons to think you may be mis-remembering things<sup>3</sup>, or lose track of the time, or lose track of who you are<sup>4</sup>? Maybe then it is fine not have confidence  $x$  in  $p$ . The principle should apply only if there are no other relevant changes to your doxastic state. Still better:

*Conditional Updating (3)*

If, prior to learning that  $q$ , your confidence in  $p$  on the supposition that  $q$  should be  $x$ , then, after learning that  $q$  and being subject to no other relevant changes to your doxastic state ( $q$  is the strongest thing you learn and you do not suffer any doxastic corrosion) your confidence in  $p$  should be  $x$ .

With the Bayesian idea made precise, we can apply it to this case. Let  $Six$  be the proposition that your die landed six. Let *Infinity* be all that you learn at Stage Two (that there are infinitely many people in a position similar to yours... etc.) Let *Central Point* be all that you learn at Stage Three (that the universe has a central point, location unknown to all dice-throwers). And let *Isa* be all that you learn at Stage Four (that your buddy's name is Isa... etc.)

At Stage One, your conditional confidence in  $Six|Infinity$  should be  $1/6$  (What does the universe being finite or infinite have to do with your die? Your conditional confidence in  $Six|Infinity$  should match your conditional confidence in  $Six|\neg Infinity$ . So, by the *No*

---

<sup>2</sup> For an early discussion of failures of *Conditional Updating (2)* due to forgetting, see Talbot (1991). For a very systematic recent effort to solve the problem see Titelbaum (2013).

<sup>3</sup> A classic example in which *Conditional Updating (2)* fails for this reason is Frank Arntzenius' 'Two Roads to Shangri La', discussed in Arntzenius (2003).

<sup>4</sup> The classic example in which *Conditional Updating (2)* fails for this reason is Adam Elga's 'Sleeping Beauty', discussed in Elga (2000).

*Difference Principle*, your conditional confidence in  $Six|Infinity$  should match your unconditional confidence in  $Six$  – at  $1/6$ ). As you transition to Stage Two you learn *Infinity*, and nothing stronger. You suffer no doxastic corrosion. So, by *Conditional Updating (3)*, at Stage Two your confidence in  $Six$  should be  $1/6$ .

At Stage Two, your conditional confidence in  $Six|Central\ Point$  should be  $1/6$  (What does the existence of a universal central point have to do with your die?) As you transition to Stage Three you learn *Central Point* and nothing stronger. You suffer no doxastic corrosion. So, by *Conditional Updating (3)*, at Stage Three your confidence in  $Six$  should be  $1/6$ .

At Stage Three, your conditional confidence in  $Six|Isa$  should be  $1/6$  (What does her being called Isa and living on a planet with two moons have to do with your die? Your buddy has to have some name – Isa is as good a name as any!) As you transition to Stage Four you learn *Isa* and nothing stronger. You suffer no doxastic corrosion. So, by *Conditional Updating (3)*, at Stage Four your confidence in  $Six$  should be  $1/6$ .

The *Symmetry* claim, on the other hand, follows naturally from the doctrine of *Time Slice Rationality*.<sup>5</sup> Let your *epistemic situation at a time* be everything you believe, know, suspect etc. at that time, plus all the evidence available to you at that time. The doctrine of Time Slice Rationality has two things to say about the relation between your epistemic situation and what you should believe. First it says that what you should believe at any given time is determined by your epistemic situation at that time. For the purposes of figuring out what you should believe today, it doesn't matter what evidence was or will be available to you yesterday or tomorrow, nor does it matter what you believed or will believe yesterday or tomorrow. All that matters is what evidence is available to you today, and what you believe today. Second it says that your beliefs, at any given time, about what epistemic situations you

---

<sup>5</sup> The basic idea is discussed and defended by Brian Hedden in Hedden (2015a) and Hedden (2015b), and by Sarah Moss in Moss (2015).

are in at other times, play no special role in determining what you should believe at that time (at least about topics that do not directly have to do with personal identity over time). For the purposes of figuring out what you should believe today (about topics that do not directly have to do with personal identity over time) it doesn't matter whether you today believe that *you* were in a certain epistemic situation yesterday, or whether you believe that *somebody else* was in that epistemic situation yesterday.

At Stage Four, you and Isa are in almost the same epistemic situation. You have all the same evidence available to you and all the same beliefs -- except that you have the belief that you would express like this: "I threw Reader's Die, and was rightly  $1/6$  confident that it landed six." She has a belief that she would express like this: "I threw Isa's Die, and was rightly  $1/6$  confident that it landed six." But the doctrine of Time Slice Rationality says that differences like this don't matter. So any confidence that you should have in *Six*, she should have in *Six*. And  $1/2$  confidence is the only fair way to go (it can't very well be that she should be  $1/6$  confident that your die is a six,  $5/6$  confident that her die is a six!)

#### **4. Reply: Dismiss the Problem**

What to think? I will consider and reject two replies, before settling on the reply I favor. The first reply is dismissive.

"When we, as philosophers, ask questions about what you should believe in this or that situation, we are in the business of building a normative epistemological theory. Why get in that business? Because we want a tool that will, first, help us figure out what to believe in situations in which it is not obvious what to believe, and, second, help us predict what other people will believe in these situations (on the assumption that other people are, most of the time, rational). To be useful the tool, like any other tool, does not have to work in every



conceivable situation. It just has to work in the situations we are likely to find ourselves in.

“This problem has to do with what to believe in a series of wild situations – situations involving infinite numbers of dice throwers communicating across every greater distances. But who cares about adapting our normative epistemological theories so that they work in these situations? Complaining of a theory that it does not work in these situations is like complaining of a satellite navigation system that it does not give directions on the surface of Venus. Neither you nor any other human has any expectation of ever going there!”

But this is too quick. Stage One is not a very wild situation. It is the sort of situation we face all the time. Normative epistemologists had better figure out what to say about Stage One. Nor is Stage Two a very wild situation. Reputable cosmologists are inclined to think that space and matter are indeed infinite<sup>6</sup>, and that events increasingly far apart are increasingly probabilistically uncorrelated. If they are right then whenever you throw a die you should have confidence 1 that there are infinitely many people like you throwing dice like it. Normative epistemologists had better figure out what to say about Stage Two. Nor is Stage Three a very wild situation. To find yourself in a situation relevantly like Stage Three all you need is to learn of the existence of a spatial point to which you and your infinite siblings can refer, but whose location is unknown to you. The first truly wild situation, such that we can say with great confidence that you will never face it, is Stage Four.

Perhaps, then, we should be satisfied with a normative epistemological theory that tells you what to believe at Stages One, Two and Three, but is silent about what to believe at

---

<sup>6</sup> Our most accurate measurements of deep space, recently performed by the Baryon Oscillation Spectroscopic Survey Collaboration, reported in Schlegel et. al (2014), indicate that the visible universe is approximately flat, approximately Euclidean. The galaxies we see when we point our telescopes north are not the galaxies we see when we point our telescopes south. In a press release accompanying the publication of that paper, David Schlegel, leader of the BOSS team, said: “One of the reasons we care is that a flat universe has implications for whether the universe is infinite. That means – while we can’t say with certainty that it will never come to an end – it’s likely the universe extends forever in space and will go on forever in time. Our results are consistent with an infinite universe.”

Stage Four. But to stop there would be to leave a pressing question unanswered: *Why* is the normative theory silent about what to believe at Stage Four? Why does a conversation with Isa bring the gears of the theory shuddering to a halt? Not because of what Isa says, surely. Much of what Isa says is predictable in advance, and the unpredictable-in-advance things she says appear to have nothing whatsoever to do with dice.

If our satellite navigation systems don't work on Venus, then we want to know *why*. Maybe we can learn something about the nature of the systems this way.

### **5. Another Reply: Reject *No Changes***

Satisfactory replies to paradoxes identify the false claim and explain why it is false yet truthy – typically by identifying a true claim in its near vicinity. Here is a shot at a satisfactory reply.

“*No Changes* is false, because your confidence should change from Stage Three to Stage Four. At Stage Three your confidence should be 1/6. At Stage Four it should be 1/2. What went wrong with the argument from conditionalization? Well, *Conditionalization (3)* is meant to apply only when you learn something and there are no other relevant changes to your epistemic situation (like learning further things, forgetting things, gaining reasons to think you may be mis-remembering things, losing track of the time, or losing track of who you are). We supposed that in this case there are no other relevant changes to your epistemic situation. That was a mistake.”

Well and good. But we need a proposal about what the *other relevant change to your epistemic situation* is. First proposal:

“At Stage Four you begin to *converse* with Isa. That changes your epistemic situation because conversation brings with it new norms on belief. When you converse with someone

who knows all you know, you should be moved only by considerations that are in principle suited to move you both. But the consideration ‘I threw Reader’s Die, and was rightly 1/6 confident that it landed six’ is not in principle suited to move Isa! You can’t very well expect her to be 1/6 confident that Reader’s Die landed six on these grounds!

“The point is that, in conversation, epistemic inquiry is a collaborative project. If you and she are moved by the considerations ‘I threw Reader’s die’ and ‘I threw Isa’s die’ respectively, then you will each regard the other as *deluded* in an especially tragic way – it is absolutely and in principle impossible to demonstrate to the other the nature of the other’s delusion. You believe everything she says. She believes everything you say. And yet you discount each other’s rightful confidence. This discounting is *uncivilized*, contrary to the norms of conversation.

“To bring out how uncivilized it is, suppose that both you and Isa have dogs, and the authorities tell you that you must together accept one of two deals:

*Deadly Reader-Six*      Both dogs die iff Reader’s die landed six, Isa’s die landed non-six.

*Deadly Isa-Six*      Both dogs die iff Isa’s die landed six, Reader’s die landed non-six.

If, when the time comes to choose a deal, you are moved by the consideration ‘I threw Reader’s die’ to have confidence 5/6 that her die, Isa’s die, landed six, and she is moved by the consideration ‘I threw Isa’s die’ to have confidence 5/6 that your die, Reader’s die, landed six, then things will get ugly. You will and should scream and shout and do everything in your power to lessen her say in the decision (big things are stake, remember). If the authorities bring you and her physically together (maybe they have powerful transportation devices, on top of their powerful telecommunication devices), then you will and should literally, physically fight with her over which deal to choose – even though you

and she have exactly the same goals, and you each believe everything the other says about their background situation. This is not how epistemically responsible people behave.”

But I am not happy with this proposal. For one thing, it leaves open the question of just what the temporal boundaries of the conversation, governed by ‘civilized’ norms, are. Your confidence in *Six* should be  $1/2$  as you talk with Isa. What should it be when you put the phone down? Should it immediately revert to  $1/6$  then? What should it be during interruptions in the conversation? Should you flicker between  $1/2$  and  $1/6$  as the conversation starts and stops?

For another thing the ‘civilized’ norms seem to me too concessive. Prior to talking with your buddy you should believe that most probably she/he is a sixer with grounds for confidence that she/he is a non-sixer. As you begin talking with her you should believe that most probably you are talking with her, of all people, because she is a sixer with grounds for confidence that she is a non-sixer. You can and should discount her confidence that she is a non-sixer, for just this reason. To fail to discount her confidence seems less a mark of *civilized respect*, more a mark of *weak pliability* – like a needy president who adapts his views to please whoever walks into his office.

So what else might be the ‘relevant change’ in your epistemic situation as you transition from Stage Three to Stage Four? Another proposal:

“At Stage Four you can think about Isa and her die in a new way. Previously you could only think about her under the description ‘My buddy’, which is shorthand for ‘the member of the group of dice throwers whose dice are anti-correlated with mine, with the same group number (group numbers given by distance from the central point) I have.’ Now you can think of her as ‘Isa’. You can think of her as ‘that person’. These are new thoughts, thoughts that were not available to you at Stage Three.

“The availability of these new thoughts changes your epistemic situation, because that her die is anti-correlated with yours is no longer part of the way in which you think about her and her die. Previously it was *baked in* to your way of thinking about her die that you were 5/6 confident it landed six. You could not think about it without thinking of it as a member of a group such that you are 5/6 confident that every member of it is a six. Now you can think about it directly, without prejudice concerning how it landed.”

This is, I think, the most interesting defense of the idea that you should switch. But I remain unhappy with it, for several reasons.

First, notice that a brief glimpse of your buddy’s die is enough to give you the ability to think about your buddy’s die directly, without prejudice concerning how it landed – now you can think about it as *that die*. So, if the proposal is right, a brief glimpse should change your levels of confidence. But that is *prima facie* absurd. To get a sense of the absurdity, imagine yourself at Stage Three given a forced choice between *Deadly Six* over *Deadly One Two*. You would (I take it) choose *Deadly Six*. Now imagine yourself being given a brief glimpse of your buddy’s die. Would you immediately switch to *Deadly One Two*? I would not switch. And I find very hard to believe that I would be making a mistake by not switching.

Second, notice that the proposed theory gets confused in other contexts. Suppose that, in addition to defining a *buddy* relation, the authorities define a *team-mate* relation. Each person has five team-mates. Each non-sixer is team-mates with five sixers. Nobody is team-mates with their buddy. (For the detail: let’s say that the first non-sixer and the second to sixth sixers are team-mates. The second non-sixer and the first and seventh to tenth sixers are team-mates. The third non-sixer and the eleventh to fifteenth sixers are team-mates... etc.) Suppose the authorities give you a look at your buddy’s die and a look at the die of each

of your five team-mates. According to the proposal, you should now spread your credences evenly between you and your buddy, and between you and your five team-mates. But you can't do that.

Third, notice that there is no account of why it is epistemically important that you be able to think about the die directly, without prejudice concerning its outcome. Why does this ability change what it is rational for you to believe? That you have the ability seems unimportant in light of the thought: *Probably I have the ability to think about Isa's die directly, without prejudice concerning its outcome, only because it is a six.*

## **6. The Reply I Favor: Stick to your Guns, Reject *Symmetry***

This is the reply I favor. Because I favor it, I will not embed it in quotes.

*Symmetry* is false. At Stage Four you should be 1/6 confident that you threw a six.

What is wrong with the argument from Time Slice Rationality? Well, the first premise of that argument is right: What you should believe at any given time is indeed determined by your epistemic situation at that time. But, contra the second premise, there is a significant difference between your and Isa's epistemic situation.

What about the idea that considerations to do with personal identity over time, like 'I threw Reader's die', have no special epistemic significance? Well, the key difference between your and Isa's epistemic situation is not that you have *that* de se belief. The key difference is that you have another de se belief.

When we attend to an event and try to form a view about how exactly it went, the most natural and obvious question to ask is: 'What sort of an event is it?' Less natural and obvious is: 'Why am I attending to this event, in particular?' But, natural or not, that question may be epistemically relevant.

When it comes to Reader's die, you know it to have been fair, and rightly have a default attitude of  $1/6$  confidence that it came up six. Your explanation for why you are attending to it now (I threw it, and I have been attending to it ever since) gives you no reason to reassess your default attitude. When it comes to Isa's die, you also know it to have been fair, but your explanation for why you are attending to it now (I threw Reader's die, and this one was selected by the authorities from a group of dice whose outcomes are anti-correlated with the outcome of my throw) does give you reason to err from the default attitude.

Notice that the fact that it was *you* who threw Reader's die plays no special role in the explanation. Suppose that one more thing happens:

#### *Stage Five*

Your mother wanders into your room. She finds you there, sitting before a table on which there is closed cigar box, talking on a green phone. You fill her in on everything that has been going on – all about the die, the infinite universe, the central point, your conversation with Isa... etc.

In this situation, though your mother certainly should not think 'I threw Reader's die', she should nonetheless be  $1/6$  confident that Reader's die is a six,  $5/6$  confident that Isa's die is a six. She knows Reader's die to have been fair, and her best explanation of why she is attending to it (something like: *because my dear child threw it and I wandered into the room*) gives her no reason to stray from  $1/6$  confidence that it landed six. She also knows Isa's die to have been fair, but her best explanation for why she is attending to it (something like: *because it is one of those dice whose outcomes are inversely correlated with the outcome of Reader's die*) does give her reason to stray from  $1/6$  confidence that it landed six.

Or, more fancifully, suppose that things go a little bit differently at Stage Four.

*Stage Four\**

Before the phone rings you are placed in a *psychology scrambler*. A person, Alex, emerges from the scrambler. Alex has your body, but Alex's present mental states bear no interesting relations to your past mental states. The phone rings. Alex picks it up and learns all about Isa, the dice, the desks, the scramblers, etc.

Is Alex you? How you answer this question will depend on whether you endorse some version of the psychological criterion of personal identity over time or some version of the bodily criterion of personal identity over time. But notice that, on our proposal, the question has no bearing on what Alex should believe. No matter which answer is correct, when Alex wonders "Why am I attending to Reader's die?", the proper explanation (If the bodily criterion is correct: *because I threw it*. If the psychological criterion is correct: *because the previous owner of this body threw it*.) gives Alex no reason to stray from a default 1/6 confidence that Reader's fair die landed six. But still when Alex wonders "Why am I attending to Isa's die?" the best explanation (still: *because it is one of those dice whose outcomes are inversely correlated with the outcome of Reader's die*) does give Alex reason to stray from a default 1/6 confidence that Isa's fair die landed six.

In sum: *de se* beliefs concerning how it came to be that I am in the situation I am in play a special role in epistemic rationality, though beliefs about personal identity over time do not.



## References

Elga, Adam (2000). 'Self-Locating Belief and the Sleeping Beauty Problem', *Analysis*, 60(2): 143-147, 2000.

Hedden, Brian (2015a). *Reasons without Persons: Rationality, Identity, and Time*. Oxford: Oxford University Press.

Hedden, Brian (2015b). 'Time-Slice Rationality'. *Mind*, 124(494), 449-491.

Moss, Sarah (2015). 'Time-Slice Epistemology and Action Under Indeterminacy', *Oxford Studies in Epistemology*, John Hawthorne and Tamar Szabó Gendler, eds. Oxford, 2015.

Schlegel, David, et al. (2014). 'The Clustering of Galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations in the Data Release 10 and 11 Galaxy Samples', *Monthly Notices of the Royal Astronomical Society* arXiv:1312.4877v2

Talbott, William. 1991. 'Two Principles of Bayesian Epistemology', *Philosophical Studies*, vol. 62: 135–50

Titelbaum, Michael. 2013. *Quitting Certainties: A Bayesian Framework for Modeling Degrees of Belief*, Oxford University Press, Oxford.