



# Bayesianism and self-doubt

Darren Bradley<sup>1</sup> 

Received: 13 February 2020 / Accepted: 16 September 2020 / Published online: 20 November 2020  
© The Author(s) 2020

## Abstract

How should we respond to evidence when our evidence indicates that we are rationally impaired? I will defend a novel answer based on the analogy between self-doubt and memory loss. To believe that one is now impaired and previously was not is to believe that one's epistemic position has deteriorated. Memory loss is also a form of epistemic deterioration. I argue that agents who suffer from epistemic deterioration should return to the priors they had at an earlier time. I develop this argument regarding memory loss then extend it to cases of self-doubt.

**Keywords** Self-doubt · Forgetting · Bayesianism

## 1 Introduction

How should we respond to evidence when our evidence indicates that we are rationally impaired? Christensen (2010) and Schoenfield (2018) have argued that there is a tension between conditionalization and the belief that you are rationally impaired. I offer a novel theory of how we should respond when we believe we are rationally impaired. I will argue that an extension of conditionalization applies in cases where the agent loses information, and then argue that cases of self-doubt can be assimilated to cases of losing information.

Section 2 explains a case of believed rational impairment and Sect. 3 describes Christensen's and Schoenfield's positions. Section 4 shows how agents can respond to memory loss, then argues that a similar response is called for when self-doubting evidence is acquired. Section 5 applies the account to a series of cases that build up to the original case of believed rational impairment. Section 6 argues that this strategy can be used to guide agents by taking a step back and focussing on the synchronic constraints needed for both memory loss and self-doubt. Section 7 discusses objections. Section 8 concludes.

---

✉ Darren Bradley  
d.j.bradley@leeds.ac.uk

<sup>1</sup> Philosophy Department, Woodhouse Lane, Leeds University, Leeds LS2 9JT, UK

## 2 Drug case

Where there is a source of information, there is room for doubt about whether it is trustworthy. We'll focus on two sources of information—memories and reason. We can doubt our memories and we can doubt that we are reasoning correctly. Let's start with Christensen's example<sup>1</sup> which is most directly a case of self-doubt about *reasoning*:

### Drug

Suppose first that I'm a rational scientist investigating some phenomenon experimentally, and suppose that, were I to get evidence E, it would give me excellent reason for confidence in G (say, because E is highly unexpected, and G would be a terrific explanation for E). And suppose it's Sunday, and I'll get the results of my experiment when I get to the lab Monday morning. In this case, it seems that these things may well be true of me:

1. I'm not highly confident that G is true.
2. I am highly confident that if I will learn E tomorrow, G is true.
3. If I get to the lab and learn that E is true, I should become highly confident that G is true.

So the confidence in G that I should adopt Monday at the lab, if I do learn E, lines up with the confidence I have Sunday that G is true on the supposition that I will learn E on Monday.

Now instead of considering just the possible experimental outcome E, let's consider a more complex bit of evidence I could acquire Monday morning. I could learn not only E, but D: that a powerful explanation-assessment-disrupting drug is slipped into my breakfast coffee on Monday. Here, it seems that a gap opens up between the two things that lined up nicely before. First consider how confident I should be that if I will learn (E&D) tomorrow, G is true. My being drugged tomorrow has no bearing on the actual evidential/explanatory connection between E and G, and no independent relevance to G. So it seems that, today, I should think that, if E is true, G is very likely true, whether or not I get drugged tomorrow morning. Thus:

4. I am highly confident that if I will learn (E&D) tomorrow, G is true.

But if I actually do learn (E & D) tomorrow, will it in fact be rational for me to become highly confident in G? It seems not—after all, if I learn D tomorrow, it will not be rational for me to trust my assessments of explanatory support. And this is true whether or not I'm actually affected by the drug. So it seems, at least at first blush, that:

5. If I go to the lab and learn (E&D), I should not become highly confident that G is true.

So it seems that the higher order evidence about my being drugged produces a mismatch between my current confidence that G is true on the supposition that I will learn certain facts, and the confidence in G that I should adopt if I actually learn those facts.

<sup>1</sup> See Schoenfield (2018, p. 690) for a similar example.

(Christensen 2010, pp. 199–200 Notation altered.)

It will be useful to abstract away from the details:

G = Hypothesis

E = Evidence which supports G

G and E are eternal i.e. non-self-locating propositions. We can add to the story that you are confident throughout in the non-self-locating proposition that a powerful explanation-assessment-disrupting drug is slipped into your breakfast coffee on Monday. And we need to add the self-locating proposition that it is *now* Monday:

D- = It is *now* Monday

Given the assumption that you are confident throughout that a powerful explanation-assessment-disrupting drug is slipped into your breakfast coffee on Monday, we can make the significance of D- more explicit:

D = It is *now* Monday, a time at which I am impaired

On Sunday you are unimpaired and disbelieve D. On Monday you believe D and so do not believe you can infer G from E.

We also need to add to the story that on Monday you don't trust your memories. Otherwise, if you did trust your memory on Monday, you could recall that on Sunday you believed that G was likely given E, and could just defer to your earlier credences. So we must 'add to the description of the [possible impairment] that it also distorts one's memories of one's previous credences' (Christensen 2010, p. 200).

Let me make five clarifications. First, you trust your reasoning ability and memories otherwise. Second, you are never in fact impaired, you just believe you are (or will be). Third, you always know what time it is. Fourth, you move from being sure that you are not impaired to believing that you are. (My analysis is not intended to apply to a case where you think that you might have been impaired all along. In such a case there would be no need for memory loss so I think that conditionalization could be applied, although I won't argue for this here (see fn. 10.) These assumptions will be in play for the whole paper. An assumption we will drop in Sect. 6 is that a single unique credence is obligatory given any particular set of evidence; call this Uniqueness.<sup>2</sup>

#### *Uniqueness*

A single unique credence is obligatory given any particular set of evidence

### 3 Christensen and Schoenfield

In this section I'll sketch Christensen's and Schoenfield's responses. They both consider the examples to be a problem for conditionalization<sup>3</sup>:

<sup>2</sup> See White (2005).

<sup>3</sup> Assume that conditionalization applies when  $P(E) > 0$  and that conditional probabilities are undefined when  $P(E) = 0$ .

Let  $P_0$  represent an agent's credences at time  $t_0$  and let  $P_1$  represent their credences at time  $t_1$ .

*Conditionalization:* If the only change to an agent's epistemic position between  $t_0$  and  $t_1$  is that they come to believe  $E$  with certainty, then  $P_1(\cdot) = P_0(\cdot|E)$

Both their analyses start as follows. (Suppose for vividness and expository ease that you are the agent.) On Sunday,  $G$  is not likely, but  $G$  is likely given  $E$ :

$$P_{\text{Sunday}}(G) = \text{Not high}$$

$$P_{\text{Sunday}}(G|E) = \text{High}$$

How likely is  $G$  given  $D$ ?<sup>4</sup> Future impairment makes no difference to how likely  $G$  is, so  $D$  is irrelevant:

$$P_{\text{Sunday}}(G|E\&D) = \text{High}$$

Therefore, if  $E\&D$  is learnt by Monday, by conditionalization:

$$P_{\text{Monday, } E\&D}(G) = \text{High}$$

Intuitively, conditionalization says that  $D$  is irrelevant—you should believe what you would have believed if you had not believed that you are impaired. Call this type of response 'Steadfastness'.<sup>5</sup>

Christensen, Schoenfield and myself agree that Steadfastness is problematic. Christensen cautiously backs off from the conclusion that conditionalization should be rejected in such cases. He backs off because of the complexities that arise when self-locating evidence is learnt, but maintains that 'we can see a contrast between updating involving higher-order beliefs and ordinary cases of updating' 2010 p. 201.

Schoenfield thinks that Christensen was wrong to back off from the conclusion that conditionalization faces a problem. She agrees with Christensen that we have to work out how to update on self-locating evidence, and so she develops a theory about how to do so. However her theory still yields the same problematic answer of Steadfastness. So she draws a distinction between *conforming* to a rule and *planning* to conform to a rule. These come apart when we might fail to follow the plan. Schoenfield argues that Steadfastness is the best rule to conform to, but that the agent is unable to reliably follow the Steadfastness rule once they believe they're impaired, so should not plan to follow it. I think the underlying intuition is that one should not plan to be Steadfast because one cannot be *guided* by the Steadfastness rule. I will assume that we are looking for rules of rationality which can guide the agent (see Sect. 6 for discussion).<sup>6</sup> Schoenfield

<sup>4</sup> Here is the problematic move of trying to conditionalize on a self-locating proposition. More on this below.

<sup>5</sup> I take 'Steadfastness' from Schoenfield. She compares it to 'Calibrationism', but this term gets used in different ways [see Schoenfield (2015), fn. 6], so I prefer 'Return-to-Priors'. 'Steadfast' is also not perfectly precise. One way to be Steadfast is to allow a mismatch between higher level and first order credences (e.g. epistemic Akrasia), letting first order credences remain unaffected by higher order doubt. Another way to be Steadfast is to rule out the possibility of genuine evidence that would justify higher order doubt. I cast doubt on these positions in Bradley (2019).

<sup>6</sup> Guidance leads to internalism and the evidentialism associated with Bayesianism over externalism about justification (Goldman 1986). See Meacham (2010) and Schwarz (2012) for discussions connecting internalism and Conditionalization, and Wedgwood (2017 ch. 7) for an argument for internalism based on guidance. See Schoenfield (2018, section 2) for some of the connections to evidentialism and accuracy.

argues that the best rule to *plan* to follow is one which keeps your later credence in G not high:

$$P_{\text{Monday, E\&D}}(G) = \text{Not high}$$

Call this *Return-to-priors*. Although intuitively correct, *Return-to-priors* still conflicts with conditionalization, which implied that  $P_{\text{Monday, E\&D}}(G) = \text{High}$ .<sup>7</sup> So we still need to know exactly what goes wrong with conditionalization, and what rule should replace it. I will offer a theory which vindicates the intuition behind *Return-to-priors*. Then in Sect. 6 I will discuss how *Return-to-priors* might be generalized.

## 4 Memory loss and self-doubt

Why does conditionalization go wrong in DRUG? A necessary condition on the rationality of conditionalizing on E between  $t_0$  and  $t_1$  is that the agent is certain at  $t_0$  that the only change to their epistemic situation between  $t_0$  and  $t_1$  is that they learn some proposition with content E and come to believe E with certainty.<sup>8</sup> Let's distinguish two possibilities in which this condition is not satisfied:

- (a) A self-locating proposition changes in truth-value.
- (b) Epistemic deterioration, which includes at least memory loss, impaired reasoning abilities, believed memory loss and believed impaired reasoning abilities.

(a) allows an agent's epistemic position to change without acquiring any evidence. The proposition that it is Sunday changes from being true to being false as time passes, thus a fixed belief that 'today is Sunday' becomes false. Facts external to the agent have changed, making a true belief false, and changing the agent's epistemic position, even if nothing is learnt (with certainty or without).<sup>9</sup>

For (b), where memory loss and self-doubt are possible for the agent, the agent will not be certain at  $t_0$  that the only change to their epistemic position is that they learn E (with certainty or without).<sup>10</sup>

Both (a) and (b) appear in DRUG, but I think the core issue concerns (b) epistemic deterioration. I agree that (a) self-locating beliefs play a role in generating the self-doubt—the agent learns she is at a time when she might be impaired—but self-location is a feature of *any* change in epistemic state (for a sufficiently introspective/reflective agent<sup>11</sup>). *Learning* E locates the agent at a time after they have learnt E. That is, when an agent learns E, they are also in a position to learn the self-locating evidence 'I am

<sup>7</sup> I focus on my positive account in this paper, and engage with Schoenfield in Bradley (2020).

<sup>8</sup> It is not always made explicit that this must be the *only* change, but it is assumed by influential arguments for conditionalization, especially Greaves and Wallace (2006). See Schoenfield (2018) and Bradley (2020) for discussion.

<sup>9</sup> Notice that a change of truth-value for any reason can have this effect, not just due to self-location.

<sup>10</sup> Neither (a) nor (b) occur in a case where the agent thinks that they might have been impaired all along, and learns that they have been, which is why I think conditionalization applies in such cases.

<sup>11</sup> For non-introspective agents, the agent might learn E without learning that she is at a time after E is learnt e.g. Williamson (2000). This supports my point that self-location is not relevant. If the agent cannot make the inference from 'E' to 'I am at a time after I have learnt E', then this reason to think self-location is relevant is removed.

at a time after learning E'. So in a sense, *all evidence is self-locating*. Agents can usually ignore the fact that they are now *at a time* after learning E because that fact is not relevant to the hypothesis at issue.<sup>12</sup> Similarly, agents can here ignore the fact that they are now *at a time* when they might have forgotten E (or are reasoning unreliably from E). What's relevant is just that they might have forgotten E.

I will first show how one form of epistemic deterioration, memory loss, can be dealt with, then argue that cases of *self-doubt* can be assimilated to cases of memory loss. So I will argue in the rest of this section that (1) memory loss can be handled by a natural extension of conditionalization and (2) self-doubt has the same consequences as memory loss.

#### 4.1 Memory loss can be handled by a natural extension of conditionalization

So, how should agents respond to memory loss?<sup>13</sup> I will argue that a natural extension of conditionalization gives an answer that is sufficient for our purposes. First we need to generalize conditionalization. Conditionalization is standardly expressed with reference to time—comparing credence at one time with a later time. But change in the *time* doesn't really matter. What really matters for changing credences is the change in *information*, where changes in information include both learning and forgetting. This leads to *I-Conditionalization*:

*I-Conditionalization*:

Let  $P_A$  represent an agent's credences at time A and let  $P_B$  represent their credences at time B:

If the only change in information between time A and time B is that E is believed with certainty at time A and E is not believed with certainty at time B,  
then  $P_A(\cdot) = P_B(\cdot|E)$ <sup>14</sup>

*I-Conditionalization* is symmetrical with respect to time—it simply relates the credences at two times, one with E believed and one without. A similar principle has been defended by numerous writers and I take *I-Conditionalization* to be in the spirit of their proposals.<sup>15</sup>

I do not intend to add to their arguments, but it might be useful to see some comments of Brian Skyrms and Brian Hedden to motivate *I-Conditionalization*. Skyrms writes:

Our system could start with a given probability distribution, and instead of continually updating, simply keep track of what it has learned. At any stage of the game its current probability distribution will be encoded as a pair whose first member is the original prior distribution, and whose second member is the total

<sup>12</sup> See Bradley (2011, p. 409) for a defence of the approach I am taking here; see Bostrom (2007) for dissent.

<sup>13</sup> This is a perennial problem for Bayesians; see Talbott (1991) and Williamson (2000, p. 219).

<sup>14</sup> To keep things simple, we can assume that information is identified with evidence with probability 1, so that E is believed with certainty. Titelbaum (2013) focusses on 'quitting certainties', but Hedden (2015, pp. 37–40) points out that 'credence 1 is not the issue' for memory loss.

<sup>15</sup> See Levi (1980), Skyrms (1983), Meacham (2008), Titelbaum (2013, section 6.1), Hedden (2015), Moss (2015), Meacham (2016) and Builes (forthcoming).

evidence to date. If it needs a current probability, it computes it by conditionalization on its total evidence, p. 159.

This system will satisfy *I-Conditionalization*, as the effect on this system of changing from not believing E with certainty to believing E with certainty is a shift from  $P(\cdot)$  to  $P(\cdot|E)$ . Chronological order doesn't matter.

Similarly, Hedden (2015) argues that your credences at any time should equal your rational priors conditional on your evidence, resulting in what he calls 'Synchronic Conditionalization':

*Synchronic conditionalization*

Let P be the uniquely rational prior probability function [assuming Uniqueness].

If at time t you have total evidence E, your credence at t in each proposition H should equal  $P(H|E)$ .

In my view this is a plausible principle which helps support *I-Conditionalization*.<sup>16</sup> (Uniqueness is invoked here but we will drop this assumption in Sect. 6.)

Moving on, from *I-Conditionalization* we can derive a rule about the rational response when there is memory loss:

*Memory-Loss Rule*

If the only change in information between an earlier time t0 and a later time t1 is that E is believed with certainty at t0 and not believed with certainty at t1,

then  $P_{t_0 E}(\cdot) = P_{t_1}(\cdot|E)$ .

This is conditionalization with the times t0 and t1 reversed. We can combine the Memory-loss rule and the familiar temporal version of conditionalization in a principle that covers three times:

*Learn-Forget Rule-*

If the only change in information between an earlier time t0, a later time t1 and an even later time t2, is that E is not believed with certainty at t0, believed with certainty at t1 and not believed with certainty at t2,

then  $P_{t_0}(\cdot|E) = P_{t_1 E}(\cdot) = P_{t_2}(\cdot|E)$ .

As written, this is only true of the credences conditional on E. Let's extend the rule to the unconditional credences by adding 'and  $P_{t_0}(\cdot) = P_{t_2}(\cdot)$ ':

<sup>16</sup> Someone might object that Synchronic Conditionalization is a synchronic principle (saying what an agent should believe at a time), I-Conditionalization is a diachronic principle (connecting credence for an agent at one time with credence at another time), and that we cannot or should not derive diachronic from synchronic principles.

In response, synchronic and diachronic principles need not be in competition. It is plausible that Synchronic Conditionalization is more fundamental, and diachronic principles like Conditionalization and I-Conditionalization can be derived from such fundamental principles given further assumptions about the identity of an agent over time and the evolution of their beliefs e.g. information retention. Specifically, an agent who conforms to Synchronic Conditionalization and never loses information will end up conforming to I-Conditionalization. Hedden (2015, p. 8) rejects diachronic constraints, but I think his arguments only show that they are non-fundamental. See Hlobil (2015) on the need for diachronic norms. I appeal to synchronic norms because I'm doubtful that diachronic arguments like Greaves and Wallace (2006) can be made to work for I-Conditionalization. If you expect to epistemically deteriorate then the way to maximize accuracy is to fix your credence where it is. (This is the rationale behind Steadfasting).

### *Learn-Forget Rule*

If the only change in information between an earlier time  $t_0$ , a later time  $t_1$  and an even later time  $t_2$ , is that  $E$  is not believed with certainty at  $t_0$ , believed with certainty at  $t_1$  and not believed with certainty at  $t_2$ ,

then  $P_{t_0}(.|E) = P_{t_1}E(.) = P_{t_2}(.|E)$  and  $P_{t_0}(.)=P_{t_2}(.)$ .

The *Learn-Forget Rule* is symmetrical—the shifts in credence look the same whether we move forward in time from  $t_0$  to  $t_2$  or back in time from  $t_2$  to  $t_0$ . The justification is that the process of learning then forgetting should leave the agent with exactly the same credence she started with.

The final complication here is that one can learn something without any credence going to 1. For example, credence in  $E$  might rise, but not all the way to 1. There is a widely accepted method of generalizing conditionalization to such cases—Jeffrey Conditionalization:<sup>17</sup>

### *Jeffrey Conditionalization*

$$P1(.) = P0(.|E) \times P1(E) + P0(.| \neg E) \times P1(\neg E).$$

We can derive a Jeffrey-version of the *Learn-Forget Rule* showing what happens when  $E$  is partially learnt and, inversely, partially forgotten:

### *Learn-Forget Rule, Jeffrey*

If the only change in information between an earlier time  $t_0$ , a later time  $t_1$  and an even later time  $t_2$ , is that the agent undergoes a learning experience between  $t_0$  and  $t_1$ , and then forgets the learning experience between  $t_1$  and  $t_2$ :

$$P_{t_0}(.|E) \times P_{t_1}(E) + P_{t_0}(.|\neg E) \times P_1(\neg E) = P_{t_1}(.)=P_{t_2}(.|E) \times P_{t_1}(E) + P_{t_2}(.|\neg E) \times P_{t_1}(\neg E)$$

and  $P_{t_0}(.)=P_{t_2}(.)$

That's the theory for memory loss, but what does this have to do with self-doubt?

## 4.2 Self-doubt has (almost) the same consequences as memory loss

When an agent doubts their reasoning from  $E$ , they should believe  $E$  while refusing to make inferences based on  $E$ . I will argue that this has almost the same result as when an agent forgets  $E$ . The results are only *almost* identical because, of course, the former believes  $E$  and the latter does not. But believing  $E$  does not have the usual effects. A normal agent who believes  $E$  will also believe (many of) the things that follow from  $E$ ; but an agent who does not draw any inferences from  $E$  will be in the odd position of having a belief in  $E$  floating free of its usual consequences. The inferential pipeline from  $E$  to its consequences has been blocked.

Let's use an example adapted from Arntzenius (2003). Suppose there are only two routes to Shangri-La—by road or sea. Compare three cases in which Frank travels to Shangri-La (assume Frank believes what he is told).

<sup>17</sup> A referee emphasizes that some hold that Jeffrey conditionalization is non-commutative i.e. experience  $E$  followed by experience  $F$  has a different result from having experience  $F$  followed by experience  $E$  (see Weisberg 2009). However, even if you think that some learning experiences are non-commutative, it remains plausible that learning then forgetting returns you to the same credences. First, learning and forgetting are directional, in that you cannot forget something before learning it. Second, forgetting is not a learning experience.

### A. Memory loss

Frank is told that on arrival *all* visitors have their memory of the journey removed. On arrival Frank has no memory of how he got there.

### B. Complete self-doubt about memory

Frank is told that on arrival visitors who came by sea will have their memory *replaced* by false memories of a journey by road, while visitors who came by road will not be tampered with.

On arrival Frank has an apparent memory of travelling by road.

### C. Complete self-doubt about reasoning

Frank is told that on arrival all visitors will have their ability to make correct *inferences* based on their memories of the journey impaired.

On arrival Frank has an apparent memory of travelling by road.

In each story Frank, on arriving, should have the same credence that he travelled by road—it should be whatever the prior probability was before he set out.

In (A) MEMORY LOSS Frank obviously has no additional information compared to what he started out with, so has no reason to diverge from his prior. I will argue that the agents in both the self-doubt cases, B and C, are in the same position.

In (B) COMPLETE SELF-DOUBT ABOUT MEMORY Frank has an apparent memory of travelling by road, but believes that every visitor to Shangri-La has the apparent memory of travelling by road—even if they travelled by sea. So this apparent memory should be ignored.<sup>18</sup> After all, the existence of this apparent memory had a probability of 1 from the start of the story, and no piece of evidence with probability 1 can confirm anything.

Similarly, in (C) COMPLETE SELF-DOUBT ABOUT REASONING, although Frank believes that he has a genuine memory, he cannot trust the inferences he makes based on the memory.<sup>19</sup> So he should refuse to draw any inferences. Although his credence that he has a genuine memory of the road journey goes up, this memory won't have its usual consequences e.g. a high credence that he travelled by road.

Let me elaborate on why B and C should be treated the same. Imagine that Frank knows that whether he travelled by road or sea was determined by the flip of a fair coin. Now we have a value for the prior, 0.5. I take it to be uncontroversial that if everyone has an apparent memory of travelling by road, as in SELF-DOUBT ABOUT MEMORY, then Frank's credence that he travelled by road should return to 0.5 on arrival. Similarly, if he doesn't trust any inferences he draws based on his apparent memory, as in SELF-DOUBT ABOUT REASONING, he should also stick with his prior value of 0.5 on arrival. He is not in a position to take his apparent memory to confirm that he travelled by road, and there are no other beliefs, (apparent) memories or experiences which he could take to confirm that he travelled by road.

And it doesn't matter whether the prior was fixed by an objective probability, an inductive probability, or even a subjective probability. How his prior was determined is

<sup>18</sup> See Arntzenius (2003, p. 156).

<sup>19</sup> Compare Elga (2004, section 2).

irrelevant; the issue at hand is whether Frank has evidence which *favours* the hypothesis that he travelled by road.<sup>20</sup> And the evidence favours neither.

For another argument that B and C should be treated the same, consider a more realistic case with a *non-obvious* inference. Suppose you start off with a prior of 0.5 that Fermat's Last Theorem is true. Being a brilliant mathematician, you spend a day proving it true. As you smugly close your notebook, one of two things happen. In one case, you are told that every mathematician who just proved a theorem false has been given the false memory of having proved it true.<sup>21</sup> In the other case, you are told that your mathematical abilities have been disrupted all day, making your conclusions no more reliable than chance. Surely in both cases you should return to your prior credence that the theorem is true. At the very least, I take it that this is plausible and that my opponent who thinks the cases should be treated differently requires a reason to think so.<sup>22</sup> (I consider a case where there might be such a reason in Sect. 6.)

Let's return to the Shangri-La story and make it more precise. We need to distinguish two propositions:

E = I have an apparent memory of a journey by road<sup>23</sup>

G = I journeyed by road

In MEMORY LOSS credence in E is not high (perhaps 0) after arrival, and as a result credence in G is not high after arrival. This is naturally described as forgetting E. In SELF-DOUBT ABOUT MEMORY credence in E remains high but credence in G falls after arrival because Frank no longer makes an inference from E to G, as he no longer trusts that the apparent memory is accurate. In SELF-DOUBT ABOUT REASONING credence in E remains high but credence in G falls after arrival because Frank no longer makes an inference from E to G, as he no longer trusts that he can make any reliable inferences from E, even the trivial inference that he came by road. The effect on G is the same, and we can use this to motivate a rule:

#### *Doubt-Forget Rule*

Complete self-doubt about the inference from E to G has the same effect on credence in G as forgetting E. (Reference to '*the inference*' presupposes that there is exactly one inferential path from E to G. If there is no inference or more than one inference from E to G then there is a presupposition failure and the rule does not apply. See Sect. 7 for relaxation of these assumptions.)

The final complication here is that agents can have partial self-doubt e.g.

<sup>20</sup> See Wilson (2014) and Bradley (2015) for relevant discussion.

<sup>21</sup> A complication is that proving it true might be much easier than proving it false, in which case you should think the theorem is true. But we can stipulate that proving it false is as epistemically likely as proving it true. We can stipulate that the case is structurally identical to Shangri-La.

<sup>22</sup> A referee suggests that evidence is defined partly in terms of its relevance to other propositions. On this account of evidence, it is straightforward why forgetting the relevance of some event to some hypothesis should be treated as a case of simply forgetting evidence: by definition, this is indeed a case of forgetting some particular representation of the evidence, E. This supports my view, but I don't want to be committed to this theory of evidence.

<sup>23</sup> We could take E be the phenomenal state of having an apparent memory of travelling by road, but it is typical to take propositions rather than phenomenal states to have logical/evidential relations to other propositions.

#### D. Partial self-doubt about reasoning

Frank is told that on arrival *some* visitors will have their ability to make correct inferences based on their memories of the journey impaired.

On arrival Frank has an apparent memory of travelling by road.

This better models real-life cases where agents aren't certain whether they are impaired. I think we can extend the above analysis to these cases by generalizing the *Doubt-Forget Rule* to the *Partial Doubt-Forget Rule*:

#### *Partial Doubt-Forget Rule*

Partial self-doubt about the inference from E to G has the same effect on credence in G as partially forgetting E.

Putting this together, partial self-doubt should be treated the same way as partial forgetting, which should be treated as the inverse of partial learning, which is modelled by Jeffrey Conditionalization. For example, consider a case where credence in E moves from 0.6 to 1, but the agent is only 50% certain that they can make correct inferences based on E. The effect on G will be the same as moving to a credence of 0.8 in E and Jeffrey conditionalizing. After all, the agent will be 50% confident that they should make the usual inference based on E, so should move half way towards doing so—and that is exactly the shift provided by Jeffrey Conditionalization with E at 0.8.

Someone might object that (C) and (D) require qualitatively different treatments i.e. complete self-doubt and partial self-doubt require qualitatively different treatments. Perhaps when there is any possibility that the agent can make reliable inferences, they should just go ahead and make those inferences.

But this would be hard to defend. We are modelling the case of complete self-doubt with the story of Frank being told that *every* visitor will have their ability to make correct inferences based on their memories of the journey impaired. We move to partial self-doubt by changing things so *not every* visitor is so impaired. Perhaps one in a million visitors will not be tampered with. It is hard to see why these cases should be treated qualitatively differently. Surely Frank would be reckless to assume that he is the one in a million who is unimpaired.

Still, the main point of the paper is to link memory loss with self-doubt, which occurs between B and C, so for the rest of the paper I will focus on this step of the argument. The parts are now in place so it remains to work through a case and discuss objections.

## 5 DRUG again

We have established two rules:

#### (1) *Learn-Forget Rule*

If the only change in information between an earlier time  $t_0$ , a later time  $t_1$  and an even later time  $t_2$ , is that E is not believed with certainty at  $t_0$ , believed with certainty at  $t_1$  and not believed with certainty at  $t_2$ ,

then:  $P_{t_0}(.|E) = P_{t_1}E(.) = P_{t_2}(.|E)$  and  $P_{t_0}(.) = P_{t_2}(.)$

(2) *Doubt-Forget Rule*<sup>24</sup>

Complete self-doubt about the inference from E to G has the same effect on credence in G as forgetting E.

We now need to apply them to DRUG. Let's start with a simpler version of DRUG where there is memory loss on the model of MEMORY LOSS.

Forgetful experimenter

On Saturday David takes E to support G. On Sunday he learns E and infers G. On Monday he forgets E.

Let's go through David's credences on each day. I'll add some plausible credences for concreteness.

*Saturday*

David has prior probabilities.

$$P(E) = 0.5$$

$$P(G|E) = 0.9$$

$$P(G|\neg E) = 0.1$$

$$P(G) = 0.5$$

*Sunday*

David learns E. There is unproblematic conditionalization on E.

$$P_E(E) = 1$$

$$P_E(G|E) = 0.9$$

$$P_E(G|\neg E) = \text{undefined (because } \neg E \text{ has 0 probability)}$$

$$P_E(G) = 0.9$$

*Monday*

David forgets E. By the *Learn-forget rule*,  $P(G)$  returns to the Saturday level.

$$P(E) = 0.5$$

$$P(G|E) = 0.9$$

$$P(G|\neg E) = 0.1$$

$$P(G) = 0.5$$

To change this case into DRUG we just need to change forgetting E into not making inferences based on E on the model of COMPLETE SELF-DOUBT ABOUT REASONING (I omit the intermediate step where the agent becomes uncertain that a memory is accurate on the model of SELF-DOUBT ABOUT MEMORY).

Drug

On Saturday David takes E to support G. On Sunday he learns E. On Monday he learns that he is impaired by a drug which prevents him from reliably making inferences from E. David knows all this.

*Saturday*

David has prior probabilities.

$$P(E) = 0.5$$

<sup>24</sup> The hyphen plays a different role in these rules. For the *Learn-Forget Rule*, the agent learns then forgets. The *Doubt-Forget Rule* allows the theorist to connect doubt with forgetting.

$$\begin{aligned}P(G|E) &= 0.9 \\P(G| \neg E) &= 0.1 \\P(G) &= 0.5\end{aligned}$$

### Sunday

David learns E. There is unproblematic conditionalization on E.

$$\begin{aligned}P_E(E) &= 1 \\P_E(G|E) &= 0.9 \\P_E(G| \neg E) &= \text{undefined} \\P_E(G) &= 0.9\end{aligned}$$

### Monday

David is certain that he cannot make reliable inferences from E. By the *Doubt-Forget Rule*, as far as making inferences goes, he may as well have forgotten E. What happens if he forgets E? By the *Learn-forget rule*,  $P(G)$  returns to the Saturday level.<sup>25</sup>

$$\begin{aligned}P_{E-}(E) &= 1 \\P_{E-}(G|E) &= 0.5 \\P_{E-}(G| \neg E) &= \text{undefined} \\P_{E-}(G) &= 0.5\end{aligned}$$

Thus, the agent ends up not believing G. David no longer believes that E supports G, so credence in G returns to the level it had before E was learnt; this is the pre-theoretic, intuitive answer (Return-to-priors).

Comparison of the tables allows us to see the differences between forgetting and self-doubt. The differences concern the top three boxes in the far right column (italicized). In Table 1 (FORGETFUL EXPERIMENTER)  $P(E)$  on Monday falls to 0.5; in Table 2 (DRUG)  $P(E)$  stays at 1, but  $P(G|E)$  falls to 0.5. Thus we can model cases of self-doubt as cases where the evidence bears directly on conditional probabilities, in this case on  $P(G|E)$ , while holding fixed  $P(E)$ . This completes my main argument.<sup>26,27</sup>

<sup>25</sup> I use  $P_{E-}$  to represent a probability function updated on E, but without an inference to G.

<sup>26</sup> The procedure is a version of Adams conditioning (Bradley 2005), according to which, in response to some events, the conditional probabilities shift while the unconditional probability stays fixed. In this case, in response to Monday's self-doubt,  $P(G|E)$  falls while  $P(E)$  stays the same. One might now wonder why we need to bother with the comparison with forgetting. Why not respond to self-doubt by Adams conditioning (Bradley 2005, p. 351)? The reason is that Adams conditioning is not a well-defined operation. Adams conditioning refers to any change in which conditional probabilities shift while the unconditional probability stays fixed, and there are many different ways this can happen. It is plausible that we do have a clear well-defined operation in the case where, say, we learn the truth of a conditional; but cases of self-doubt merely suggest that the earlier connection from evidence to hypothesis should be discounted. But by how much? What should the new conditional probability be? This is where I suggest we appeal to credences at an earlier time, just like in the case of memory loss. Specifically, we obtained the Monday credences from the Saturday credences using the Doubt-Forget and Learn-Forget rules.

<sup>27</sup> Carnap's (1950, 1971) attempts are widely judged to have failed and the project has never recovered.

**Table 1** Forgetful experimenter

	Saturday	Sunday	Monday
P(E)	0.5	1	0.5
P(G E)	0.9	0.9	0.9
P(G ¬E)	0.1	Undefined	0.1
P(G)	0.5	0.9	0.5

**Table 2** Drug

	Saturday	Sunday	Monday
P(E)	0.5	1	1
P(G E)	0.9	0.9	0.5
P(G ¬E)	0.1	Undefined	Undefined
P(G)	0.5	0.9	0.5

## 6 Guidance and uniqueness

### 6.1 Guidance and expected priors

Someone might be concerned about whether *Return-to-Priors* could be carried out by an agent. After all, we are discussing an agent who believes that they are rationally impaired, so why should they trust any method for updating their beliefs? More specifically, we are assuming that the agent believes that their memory is impaired, so does not trust that they remember their priors, so how could they base their current credence on their priors? We saw that the problem with Steadfasting (i.e.  $P_{\text{Monday, E\&D}}(G) = \text{High}$ ) was that it cannot be followed. Similarly, how can *Return-to-Priors* guide the agent if they cannot remember what their priors were?

A few points are worth emphasizing here. First, we are not discussing agents who are *actually* impaired, just agents who believe that they are. Second, they only believe they are impaired about the inference from E to G, not that they are *generally* impaired. Third, there is no threat to the analogy between loss of memory and self-doubt. The same worries regarding guidance by *Return-to-Priors* apply to the *Learn-Forget Rule* i.e. if they lose their memory an agent cannot be guided by a rule telling them to believe whatever they believed before the memory loss.

With that said, I think *Return-to-Priors* can guide rational agents given plausible assumptions (similar points apply to the *Learn-Forget Rule*). For *Return-to-Priors* to guide rational agents, agents must be able to come to a conclusion about their priors. They cannot rely on their memory, but consider how agents arrive at priors in the first place. This is a sore point for Bayesians, as no plausible theory has been developed,<sup>28</sup> but anyone who thinks that there are rational constraints on priors will think that there must be *some* method by which rational agents can arrive at their priors. Let's call that method M. Agents can rationally believe that they are impaired about inferring from E, but unimpaired when it comes to using M to discover the prior of G. So they will rationally believe that they can use M to discover what their priors were, and

<sup>28</sup> See Meacham (2016) for discussion.

so rationally believe that they can follow *Return-to-Priors*. And as they are in fact unimpaired regarding M, their use of M will successfully lead them back to the priors they started with.

Things get more complicated if we drop the assumption of Uniqueness. Suppose that rather than there being a single rational prior credence function, there are many.<sup>29</sup> And suppose the rational prior credence function at the earlier time (before the believed impairment), call it C1, is different to the later rational prior credence function, call it C2. Then there might be a problem if the later agent does not have access to the earlier priors.

Let's run through the options. *Return-to-Priors* could be understood as recommending the C1 value of G, or as recommending the C2 value of G. There are two cases where *Return-to-Priors* can provide guidance and one where it cannot. First, if *Return-to-Priors* recommends the C1 value of G and later agents have access to C1 then they will be able to follow the rule. Second, if *Return-to-Priors* recommends the C2 value of G and the agent has access to the C2 value of G then they will be able to follow the rule. Third, coming to the problem case, if *Return-to-Priors* recommends the C1 value of G and later agents have no access to the C1 value of G then the rule cannot be used for guidance.<sup>30</sup>

Perhaps there is still hope. Even if *Return-to-Priors* recommends C1 and later agents have no access to C1, they would still have access to their current beliefs about C1. Using current beliefs about priors is suggested by Meacham (2010, section 3.2) and Schwarz (2012, section 4) in their search for a synchronic surrogate for conditionalization. So perhaps when there is self-doubt, instead of returning to the prior of G, credence in G shifts to the weighted average of the possible priors of G. Where C<sub>i</sub> ranges over the possible values of C1, the later credence in G should be:

$$C2(G) = \sum C2(C1 = C_i) * C_i(G)$$

So even an agent with no access to what the earlier priors were could follow an extension of *Return-to-Priors* with weighted possible prior credence functions. Similar points apply to agents who lose their memory, so the analogy between memory loss and self-doubt holds up to this point.

## 6.2 Guidance and expected inferences

But this leads us to a disanalogy between forgetting and self-doubt.<sup>31</sup> In cases of forgetting, any inferences that might have been based on E are irrelevant; but with self-doubt the agent can still speculate about *what inferences there might have been*. This suggests that your credence in G should be the weighted average of what you believe your credence in G would have been were you to have properly updated on E.

<sup>29</sup> Whether agents can know that they are following the rule depends on their access to the facts about which priors are rational. But I take this question to be of secondary importance; whether they are following the rule is more important than whether they know they are following the rule.

<sup>30</sup> The disanalogy is between forgetting and both self-doubt about memory and self-doubt about reasoning.

<sup>31</sup> For the related claim that we should think of Conditionalization as an ideal to aim at, see Meacham (2015).

Let me fill this out. We assumed above that when you have self-doubt you have no idea what you would have inferred from E, so you can do no better than return to the prior. But this is not a necessary part of self-doubt. You might doubt that you can make inferences from E, but still assign a high probability to E making G likely i.e. you might think it likely that the correct credence function assigns a high conditional probability of G given E. Or you might think it likely that the correct credence function assigns a low conditional probability of G given E.

Furthermore, E might even make some prior credence functions more likely than others. That is, if we assume that the agent does not have access to the priors in the usual way (method M) and has to come to a value based on weighted possible priors, then E might affect the weightings. So *Return-to-Priors* can be seen as a special case of a more general rule which makes allowances for expected priors and expected inferences:

$$C2_E(G) = \sum C2_E(C1 = C_i) * C_i(G|E)$$

This is the situation where B and C might be treated differently, alluded to in Sect. 4. Perhaps Frank thinks that he is more likely to have an apparent memory of a journey by road if he did travel by road, despite not being able to make an inference from his memory of a journey by road.

## 7 Objections

Before concluding let's discuss two objections.

### 7.1 Disanalogy: two inferences

There is a further disanalogy between forgetting and self-doubt. In cases of forgetting, *all* inferences based on E are lost; but it is easy to imagine self-doubt about complex inferences based on E without self-doubt about simple inferences based on E. How should agents respond in such cases?

Recall that the *Doubt-Forget Rule* makes reference to an inference from E to G:

#### *Doubt-Forget Rule*

Self-doubt about the inference from E to G has the same effect on credence in G as forgetting E.

Assume there is self-doubt about complex inferences based on E without self-doubt about simple inferences based on E. Where there is a simple inference from E to G, there will be no self-doubt about the inference, so the *Doubt-Forget Rule* will not apply. When there is a complex inference from E to G, there will be self-doubt about the inference, so the *Doubt-Forget Rule* will apply, saying that the effect on G is the same as the effect on G of forgetting E.

There is still a problem when there is both a simple and a complex inference from E to G. Perhaps there is a short proof and a long proof of G from E. Then 'the inference' fails to refer to a unique inference and there is presupposition failure. So what should we say?

One option would be to stipulate that there is only one inferential path from E to G e.g.

*Doubt-Forget Rule+*

When there is only one inferential path from E to G, self-doubt about the inference from E to G has the same effect on credence in G as forgetting E.

But this doesn't address the case where there are two inferential paths. And here there is an important difference between forgetting and self-doubt. Self-doubt about the complex inference from E to G does not have the same effect on G as forgetting E—forgetting E would reduce credence in G while self-doubt about the complex inference would not reduce credence in E (as G is independently supported by the simple proof based on E). A more sophisticated rule is needed. I think it's clear what the desired outcome is—credence in G should remain the same (due to the simple inference), but is less robust, as now the credence in G is resting on only the short proof rather than both the short proof and the long proof. I will not attempt to express this as a pithy principle.

## 7.2 Probabilism

We can imagine that the inference from E to G is entirely mathematical. This means that in doubting the inference the agent will violate an axiom of the probability calculus, namely that necessary propositions get a probability of 1. Violating probabilism is incompatible with standard Bayesian modelling, thus undermining the Bayesian framework we have been using.

We can compare this to the problem of logical omniscience. The problem of logical omniscience is that 'the assumption that degrees of belief satisfy the probability laws implies omniscience about deductive logic... This seems to be an unrealistic standard for human beings' (Talbot 2016).<sup>32</sup>

Our problem is different. Our problem is that on our account even *ideal* agents violate the laws of probability.<sup>33</sup> It's not that probabilism is too demanding for us mortals; it's that probabilism fails to describe ideal agents.<sup>34</sup>

Still, the problem of logical omniscience means that we already have to model agents who violate logical laws, and there are various suggestions for how to do this which we can borrow. For example, Elga and Rayo (forthcoming), following Stalnaker (1991), suggest modelling non-probabilistic agents as fragmented agents, with more than one credence function. A second option is to assume that agents are probabilistic for a given number of steps of inference, and then non-probabilistic for propositions that require extra steps of reasoning (Jago 2014). We could extend this to say that the

<sup>32</sup> Interestingly, our problem alleviates the problem of logical omniscience. If even ideal agents fail to be probabilistic, then probabilism is not as central to ideality as we might have thought. So perhaps our failure to be probabilistic is not such a damning indictment of us. Perhaps our divergence from probabilism is of a kind with ideal agents' divergence from probabilism.

<sup>33</sup> If we are to describe agents who violate probabilism, can we use the language of credences? I think so. I am even told by a referee that 'credences' was used partly to avoid using 'probabilities'. See Dogramaci (2018) for a non-probabilistic use of 'credence'.

<sup>34</sup> See Christensen (2007) for a sympathetic discussion.

number of steps for which the agent is probabilistic depends on the topic, so agents might not be able to make good inferences from E, while being able to make good inferences from other propositions. A third option is to posit mental states other than beliefs (allowing us to characterize the agent's "doxastic possibilities") in order to model rational illogical agents (Williams 2018). Of course, there are complications working out the details, but the main point is that the problem is not a specific problem for our theory.

If ideal agents are not probabilistic then the standard probabilistic framework should be thought of as only approximating what ideal agents would believe rather than perfectly representing them.<sup>35</sup> This is made more plausible when we distinguish normative ideals from idealizations.<sup>36</sup> Bayesian models, like most models, make idealizations—that is, they make assumptions that are known to be false in order to make the model easier to work with. Probabilism is such an assumption. So probabilistic agents in Bayesian models are idealized; but they are not necessarily normatively ideal, in the sense of being perfectly rational i.e. as they ought to be.

## 8 Conclusion

To believe that one is now impaired and previously was not is to believe that one's epistemic position has deteriorated. Memory loss and self-doubt are two forms of epistemic deterioration. I proposed a theory of how agents should respond to memory loss—credence in propositions based on the memory should revert to the prior. And with self-doubt being another form of epistemic deterioration, the rational response to self-doubt about reasoning is also to revert to the prior. Assuming Uniqueness, the prior will be accessible after these kinds of deterioration. If the prior is not accessible, you should revert to your expected prior, and perhaps to your expected prior conditional on what you expect the correct posterior to be given your evidence. I argued that this theory can guide agents, but it suggests that even ideal agents should be represented with non-probabilistic models.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

Bostrom, N. (2007). Sleeping beauty and self-location: A hybrid model. *Synthese*, 157(1), 59–78.

<sup>35</sup> See Hájek (2006).

<sup>36</sup> Thanks to Adam Elga, Ed Elliott, Maria Lasonen-Aarnio, Miriam Schoenfield, Robbie Williams, Alex Worsnip and several referees for helpful feedback. This research was partly funded by British Academy Mid-Career Fellowship MD19\190046.

- Bradley, R. (2005). Radical probabilism and bayesian conditioning. *Philosophy of Science*, 72(2), 342–364.
- Bradley, D. (2011). Self-location is no problem for conditionalization. *Synthese*, 182(3), 393–411.
- Bradley, D. (2019). Are there indefeasible epistemic rules? *Philosopher's Imprint*, 19(3), 1–19.
- Bradley, D. (2020). Self-locating belief and updating on learning. *Mind*, 129(514), 579–584. <https://doi.org/10.1093/mind/fzy046>.
- Builes, D. (forthcoming). Time-slice rationality and self-locating belief. *Philosophical Studies*.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81(1), 185–215.
- Elga, A. (2004). Defeating Dr Evil with self-locating belief. *Philosophy and Phenomenological Research*, 69(2), 383–396.
- Elga, A., & Agustin, R. (Forthcoming). Fragmentation and information access. In C. Borgoni, D. Kindermann, & A. Onofri (Eds.), *The fragmented mind*. Oxford: Oxford University Press.
- Greaves, H., & Wallace, D. (2006). Justifying conditionalisation: conditionalisation maximizes expected epistemic utility. *Mind*, 115(459), 607–632.
- Hájek, A. (2006). *Masses of formal philosophy 'interview'*. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.124.5339&rep=rep1&type=pdf>.
- Hedden, B. (2015). *Reasons without persons: Rationality, identity, and time*. Oxford: Oxford University Press.
- Hlobil, U. (2015). There are diachronic norms of rationality. *Thought: A Journal of Philosophy*, 4(1), 38–45.
- Meacham, Christopher J. G. (2008). Sleeping beauty and the dynamics of de se beliefs. *Philosophical Studies*, 138(2), 245–269.
- Meacham, C. J. G. (2010). Unravelling the tangled web: Continuity, internalism, non-uniqueness and self-locating beliefs. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 3, pp. 86–125). Oxford: OUP.
- Meacham, C. J. G. (2015). Understanding conditionalization. *Canadian Journal of Philosophy*, 45(5), 767–797.
- Meacham, C. J. G. (2016). Ur-priors, conditionalization, and Ur-prior conditionalization. *Ergo: An Open Access Journal of Philosophy*. <https://doi.org/10.3998/ergo.12405314.0003.017>.
- Moss, S. (2015). Time-slice epistemology and action under indeterminacy. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 172–194). Oxford: Oxford University Press.
- Schoenfield, M. (2015). A dilemma for calibrationism. *Philosophy and Phenomenological Research*, 91(2), 425–455.
- Schoenfield, M. (2018). An accuracy based approach to higher order evidence. *Philosophy and Phenomenological Research*, 96, 690–715. <https://doi.org/10.1111/phpr.12329>.
- Schwarz, W. (2012). Changing minds in a changing world. *Philosophical Studies*, 159(2), 219–239. <https://doi.org/10.1007/s11098-011-9699-0>.
- Stalnaker, R. (1991). The problem of logical omniscience. *Synthese*, 89(3), 425–440.
- Talbott, W. (1991). Two principles of bayesian epistemology. *Philosophical Studies*, 62(2), 135–150.
- Talbott, W. (2016). Bayesian Epistemology. In: E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*, Winter 2016 edition. <https://plato.stanford.edu/archives/win2016/entries/epistemology-bayesian>.
- Titelbaum, M. G. (2013). *Quitting certainties: A Bayesian framework modeling degrees of belief*. Oxford: Oxford University Press.
- Wedgwood, R. (2017). *The value of rationality*. Oxford: Oxford University Press.
- Weisberg, J. (2009). Commutativity or holism? A dilemma for conditionalizers. *British Journal for the Philosophy of Science*, 60(4), 793–812.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19, 445–459.
- Williams, J. R., & Robert, G. (2018). Rational illogicality. *Australasian Journal of Philosophy*, 96(1), 127–141.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.