

Evidential Probabilities and Credences

Forthcoming in *The British Journal for the Philosophy of Science*

Anna-Maria Asunta Eder
eder.anna-maria@uni-koeln.de

Abstract

Enjoying great popularity in decision theory, epistemology, and philosophy of science, Bayesianism as understood here is fundamentally concerned with epistemically ideal rationality. It assumes a tight connection between evidential probability and ideally rational credence, and usually interprets evidential probability in terms of such credence. Timothy Williamson challenges Bayesianism by arguing that evidential probabilities cannot be adequately interpreted as the credences of an ideal agent. From this and his assumption that evidential probabilities cannot be interpreted as the actual credences of human agents either, he concludes that no interpretation of evidential probabilities in terms of credence is adequate. I argue to the contrary. My overarching aim is to show on behalf of Bayesians how one can still interpret evidential probabilities in terms of ideally rational credence and how one can maintain a tight connection between evidential probabilities and ideally rational credence even if the former cannot be interpreted in terms of the latter. By achieving this aim I illuminate the limits and prospects of Bayesianism.

Keywords Bayesianism, Ideally Rational Credences, Epistemically Ideal Agents

1 Introduction

Enjoying great popularity in decision theory, epistemology, and philosophy of science, Bayesianism—as understood here—is fundamentally concerned with epistemically ideal rationality.¹ It assumes a tight connection between evidential probability and ideally rational credence, and usually interprets evidential probability in terms of such credence. Roughly speaking, the evidential probability of a proposition on some body of evidence captures how strongly the body of evidence supports the proposition. It captures the plausibility of the proposition given the body of evidence. Interpretations of evidential probabilities that are in terms of ideally rational credence are attractive because they establish a link between evidential relations and ideally rational credences. They thereby provide us with some reference points for our judgments of the epistemic status of our credences. Such interpretations might even have the potential to provide some regulation for the formation of our credences.

Ideal rational credences are commonly understood as the credences of an (epistemically) ideal agent. Timothy Williamson ([2002], pp. 209–11) challenges Bayesianism by arguing that evidential probabilities on some body of evidence cannot be adequately interpreted as

¹Since I am exclusively concerned with epistemic rationality and epistemically ideally rational credence, I mostly skip the reference to the epistemic dimension.

the credences of an ideal agent with that body of evidence. From this, plus his assumption that evidential probabilities on some body of evidence cannot be interpreted as the actual credences of human agents with the body of evidence either, he concludes that no interpretation of evidential probabilities in terms of credence (credence interpretation, for short) is adequate. According to Williamson, evidential probabilities are objective and independent of credences of any kind, and are understood as measuring ‘something like the intrinsic plausibility of hypotheses prior to investigation’ ([2002], p. 211).²

Despite—or perhaps because of—the possible ramifications which Williamson’s criticism of credence interpretations has for Bayesianism, it has been widely ignored. I intend to make up for this; but I do not aim to show that Williamson’s own interpretation is inadequate. Williamson’s argument is not only crucial for his own theory of knowledge but challenges Bayesianism in general. It is, thus, also relevant for decision theory, epistemology, and philosophy of science. My overarching aim is to show on behalf of Bayesians how one can still interpret evidential probabilities in terms of ideally rational credence, and how one can maintain a tight connection between evidential probabilities and ideally rational credence, even if the former cannot be interpreted in terms of the latter. By achieving this aim I illuminate the limits and prospects of Bayesianism.

I do not argue for a specific option for defending a tight connection between evidential probability and ideally rational credence. I present several options that are available to Bayesians of different tastes. I think that there is more than one adequate variant of Bayesianism. What the right option and the right Bayesian variant is depends on the purpose for which the variant is being proposed. Due to space restrictions, the focus of this paper must be kept narrow, and a discussion of the various possible purposes would exceed its scope.

I focus on Williamson’s argument against interpreting evidential probabilities on a body of evidence as the credences of an ideal agent with the body of evidence, as well as on whether the argument establishes that no credence interpretation is tenable. In Section 2, I present my reconstruction of his argument. I argue that Williamson’s argument is in need of further motivation and is, therefore, problematic. In Section 3, I show that even if the argument were not problematic, its conclusion does not establish that there are no adequate credence interpretations. In fact, I propose a credence interpretation which avoids invoking ideal agents altogether. I interpret evidential probabilities in terms of ideally rational credences that are understood in explicitly normative terms—in terms of ‘credences one ought to have’—as opposed to in terms of credences of an ideal agent. By assuming such an interpretation, on behalf of Bayesians I can present two ways to avoid Williamson’s criticism of credence interpretations. However, they come at a price that not all Bayesians might be willing to accept. In Section 4, I present a third way to uphold Bayesianism in the face of Williamson’s criticism. This way out establishes a tight link between evidential probabilities and ideally rational credences without interpreting evidential probabilities in terms of credences. It promises to be tenable for critics of credence interpretations. I conclude in Section 5 by summarizing my results.

²Patrick Maher’s ([2006]) conception of inductive probability is very similar to Williamson’s conception of evidential probability.

2 The Ideal-Agent-Credence Interpretation

2.1 The Interpretation

Williamson is right that evidential probabilities (on a body of evidence) cannot be interpreted as the actual credences of human agents (with the body of evidence)—this claim is uncontroversial and well-known.³ Thus, I focus on his argument against interpreting evidential probabilities (on a body of evidence) as credences of an ideal agent (with the body of evidence). Such an interpretation is popular among Bayesians who understand ideally rational credences in terms of credences of an ideal agent.⁴

Suppose we are speaking of a specific ideal agent s_I with a credence function Cr_{s_I} .⁵ Then, according to the ideal-agent-credence interpretation, the following holds:

(IAC) The evidential probability of a proposition p on evidence e for an agent s , $\text{Pr}_s(p|e)$, equals r if and only if it is necessarily the case that (the credence the ideal agent s_I assigns to p , $\text{Cr}_{s_I}(p)$, equals r if e equals s_I 's total evidence, $\text{tev}(s_I)$).⁶

Formally: $\text{Pr}_s(p|e) = r \leftrightarrow \Box(e = \text{tev}(s_I) \rightarrow \text{Cr}_{s_I}(p) = r)$ ^{7,8}

³This is also emphasized by Alan Hájek ([2012], Sect. 3.3.1). For an interpretation of probabilities in descriptive terms that refers to credences of us human agents, see (Bradley [2015], Ch. 5).

⁴Audrey Yap ([2014]) presents a very instructive discussion of idealizations in epistemology and the use of ideal agents.

⁵Note that ' s_I ' is a constant.

⁶I often use such brackets to make the logical structure clear.

⁷One might be inclined to replace ' $\Box(e = \text{tev}(s_I) \rightarrow \text{Cr}_{s_I}(p) = r)$ ' by ' $\Box(\text{Cr}_{s_I}(p|e) = r)$ '. However, I prefer the conditional because it gives a more detailed picture of what is required of the agent. For instance, it is required that if e is the ideal agent's total evidence, then $\text{Cr}_{s_I}(e) = 1$, to replace ' $\text{Cr}_{s_I}(e) = 1$ ' by ' $\text{Cr}_{s_I}(e|e) = 1$ ' would not reflect what is required in this case because $\text{Cr}_{s_I}(e|e) = 1$ holds independently of what the agent's total evidence is.

⁸Instead of ideally rational credences, one could interpret evidential probabilities in terms of reasoning commitments that are ideally rational, which then together with the available evidence yield credences that are ideally rational. The ideally rational reasoning commitment of an agent s with respect to a proposition p on evidence e captures s 's ideally rational judgement of the confirmational import e provides for p . It captures how s ideally rationally judges the strength of the support for p that is provided by e . s 's ideally rational credence in p then equals s 's ideally rational reasoning commitment with respect to p on e , when s possesses e as total evidence. (The reasoning commitments are in many respects like priors.) For such an account of reasoning commitments, see (Brössel [2012]), (Brössel and Eder [2014]), and (Eder and Brössel [forthcoming]). Our more fine-grained framework for modelling ideally rational epistemic states traces back to (Levi [1980], Ch. 4). However, some components of the framework are understood in slightly different ways by Isaac Levi, and Peter Brössel and me. (Following (Levi [1980]), (Brössel and Eder [2014]) refers to reasoning commitments as 'confirmational commitments'.) I phrase the discussion here in terms of credences because most Bayesians only refer to credences and because the challenge for Bayesians addressed in this paper is more intelligibly presented in terms of credences. However, I think that what I argue for can be analogously presented in terms of reasoning commitments. (In a commentary to this paper Maria Lasonen-Aarnio presents an interpretation in terms of ideal agents that is in a related spirit.)

For various reasons I usually prefer a more fine-grained way to address agents' epistemic states that distinguishes between reasoning commitments and credences. I find the following two reasons very compelling. First, Brössel ([2012]) shows that by representing epistemic states in this way one can provide an ingenious solution to the problem of old evidence in confirmation theory. Second, Brössel and I ([2014]) suggest an attractive way to resolve doxastic disagreement by introducing an aggregation rule that is applied to the reasoning commitments of disagreeing agents instead of to their credences. This way, many requirements in the literature on how to rationally agree after disagreeing can be satisfied, which are not satisfied by competing rules that are applied to credences instead.

Before I introduce my reconstruction of Williamson’s argument against interpreting evidential probabilities as the credences of the ideal agent s_I , let me clarify some things. First, Williamson presents the interpretation in an informal way in terms of a counterfactual, instead of—as it is presented here—in terms of a material conditional in the scope of a necessity operator. The truth conditions for counterfactuals are commonly given in terms of closest possible worlds, which traces back to David Lewis ([1973]) and Robert Stalnaker ([1968]). However, the present discussion is not about the credences of an ideal agent in the closest possible worlds but about her credences in all possible worlds. Thus, the phrase ‘it is necessarily the case that’, or ‘ \Box ’, is added because what is in the scope of the operator is true for ideal agents *qua* being ideal.

Second, the fact that the left side of the biconditional refers to some agent s while the right side refers to the specific ideal agent s_I might be confusing. Due to this fact the evidential probability of a proposition on some piece of evidence is the same for any arbitrary agent. It is the same because it is interpreted as the credence s_I has, given the evidence equals s_I ’s total evidence. Accordingly, there is exactly one probability function that can be interpreted in this way. This is in line with Williamson, who—I think—has already an interpretation in mind that assumes evidential probabilities to be uniquely determined, objective, and to not vary from agent to agent. Alternatively, however, one can replace ‘ s_I ’ by ‘ s_S ’, where ‘ s_S ’ refers to s ’s ideal counterpart, and thereby assume that the ideal counterpart of an agent is in all aspects like the agent except for those properties she has in virtue of being an ideal agent. Such a replacement leaves open whether evidential probabilities are uniquely determined, objective, and vary from agent to agent. What I argue in the following would analogously apply for such a replacement.

2.2 The Argument

In what follows, I present what I consider to be the best reconstruction of Williamson’s brief, informal, and ingenious argument against the ideal-agent-credence interpretation. For the sake of simplicity, I refer to the reconstruction as ‘Williamson’s argument’.⁹

The argument’s starting point is given in the following quotation:

[...] let a be a logical truth (a proposition expressed by a logically true sentence) such that in this imperfect world it is very probable on our evidence that no one has great credence in a . [...] Let b be the hypothesis that no one has great credence in a . By assumption, b is very probable on our evidence [e^*]. ([2002], pp. 209–10; notation adjusted)

In accordance with Williamson’s proposal and where a is a complex logical truth, b is as given in the quotation (that is, that no one has great credence in a), and e^* is our evidence, which we possess as human agents, the first premise runs as follows:¹⁰

⁹There might be ways to simplify the argument. However, in order to avoid further discussion about the reconstruction of the original argument I opt for the present account, which is more detailed, and clearly remains very close to the original argument.

¹⁰Note that ‘ a ’, ‘ b ’, and ‘ e^* ’ are constants.

(1^{IAC}) The evidential probability of proposition b on our evidence e^* for an agent s , $\Pr_s(b|e^*)$, is high (that is, equal to or above the appropriate threshold t , where $t > .5$).

Formally: $\Pr_s(b|e^*) \geq t$

This premise seems plausible. Our evidence e^* might support that no one has great credence in the complex logical truth a .

With this premise established, Williamson proceeds to argue that the ideal-agent-credence interpretation leads to a conclusion that is unacceptable to advocates of the interpretation, if we also make further assumptions about the credences of the ideal agent. One such assumption is reflected in the following second premise (Williamson [2002], p. 210):

(2^{IAC}) If a proposition p and a proposition q are logically equivalent, then it is necessarily the case that the credence the ideal agent s_I assigns to p , $\text{Cr}_{s_I}(p)$, equals the credence s_I assigns to q , $\text{Cr}_{s_I}(q)$.

Formally: $p \equiv q \rightarrow \Box(\text{Cr}_{s_I}(p) = \text{Cr}_{s_I}(q))$

Most epistemologists agree that this is characteristic for ideal agents *qua* being ideal. They agree that an ideal agent assigns the same credence to logically equivalent propositions.

Finally, Williamson assumes that it is characteristic for any ideal agent that she does not have great credence in propositions that are of a Moore-paradoxical form. Accordingly, where a is the logical truth referred to in the quotation above and b the proposition that no one has great credence in a , Williamson claims:

$a \wedge b$ is of the Moore-paradoxical form ‘ A and no one has great credence in the proposition that A ’; to have great credence in $a \wedge b$ would therefore be self-defeating and irrational. ([2002], p. 210; notation adjusted)

We are led to the following third premise:

(3^{IAC}) It is necessarily the case that, no matter what the ideal agent s_I ’s total evidence, $\text{tev}(s_I)$, is, the credence s_I assigns to the proposition $a \wedge b$, $\text{Cr}_{s_I}(a \wedge b)$, is not high.

Formally: $\Box(e = \text{tev}(s_I) \rightarrow \text{Cr}_{s_I}(a \wedge b) < t)$

(IAC) and the three premises are all that is required to derive the unacceptable conclusion. (IAC) and (1^{IAC}) imply:

(4^{IAC}) It is necessarily the case that the credence the ideal agent s_I assigns to b , $\text{Cr}_{s_I}(b)$, is high if e^* equals s_I ’s total evidence, $\text{tev}(s_I)$.

Formally: $\Box(e^* = \text{tev}(s_I) \rightarrow \text{Cr}_{s_I}(b) \geq t)$

Since a is a logical truth, b and $a \wedge b$ are logically equivalent. Given that they are logically equivalent, the following is implied by (2^{IAC}) together with (4^{IAC}):

(5^{IAC}) It is necessarily the case that the credence the ideal agent s_I assigns to $a \wedge b$, $\text{Cr}_{s_I}(a \wedge b)$, is high if e^* equals s_I ’s total evidence, $\text{tev}(s_I)$.

Formally: $\Box(e^* = \text{tev}(s_I) \rightarrow \text{Cr}_{s_I}(a \wedge b) \geq t)$

(3^{IAC}) and (5^{IAC}) imply:

(6^{IAC}) It is necessarily the case that the credence the ideal agent s_I assigns to $a \wedge b$, $\text{Cr}_{s_I}(a \wedge b)$, is high, as well as not high if e^* equals s_I 's total evidence, $\text{tev}(s_I)$.

Formally: $\Box(e^* = \text{tev}(s_I) \rightarrow (\text{Cr}_{s_I}(a \wedge b) \geq t \wedge \text{Cr}_{s_I}(a \wedge b) < t))$

According to (IAC), evidential probabilities are interpreted as the credences had by the ideal agent s_I with our evidence e^* . However, according to (6^{IAC}) it is necessarily the case that if the ideal agent's total evidence, $\text{tev}(s_I)$, equals our evidence, e^* , then s_I 's credence in the proposition of the Moore-paradoxical form, $a \wedge b$, is not only equal to or above the threshold, t , but also below it, which is contradictory. Thus, it is impossible that the ideal agent, s_I , has our evidence, e^* , as her total evidence, $\text{tev}(s_I)$. That is, (6^{IAC}) leads to—what I consider to be—the conclusion of Williamson's argument:

(\therefore ^{IAC}) It is not possible that our evidence e^* equals the ideal agent s_I 's total evidence.

Formally: $\neg \Diamond(e^* = \text{tev}(s_I))$

(\therefore ^{IAC}) is in accordance with the following quotation by Williamson:

Presumably, a perfectly rational being must give great credence to a , be aware of doing so, and therefore give little credence to b and so to $a \wedge b$; but then its evidence about its own states would be different from ours. If so, the hypothesis of a perfectly rational being with our evidence is impossible. There is no such thing as the credence which a perfectly rational being with our evidence would have in a given proposition. ([2002], p. 210; notation adjusted)

(IAC) seems misguided. It presupposes that for any kind of evidence it makes sense to speak about ideal agents with exactly this evidence. Thus, it presupposes that an ideal agent can have our evidence that we possess as non-ideal, human agents. (\therefore ^{IAC}) is in conflict with this, which is due to the fact that ideal agents have evidence about their own ideal doxastic state, which we do not have. Based on this, Williamson rejects the ideal-agent-credence interpretation.¹¹ For him it suffices to understand evidential probability as uniquely determined probabilities that do not represent credences but intrinsic plausibility, and measuring 'something like the intrinsic plausibility of hypotheses prior to investigation' (Williamson [2002], 211). They are supposed to reflect evidential relations that are objectively given. Williamson does not present a fully fleshed-out account of such probabilities. He claims the following in this respect:

Sometimes the best policy is to go ahead and theorize with a vague but powerful notion. One's original intuitive understanding becomes refined as a result, although rarely to the point of a definition in precise pretheoretic terms. That

¹¹In order to save credence interpretations of evidential probabilities one might be tempted to suggest a way out by restricting evidential probabilities to the total evidence that an ideal agent can have. Williamson's argument does not concern such interpretations. However, Williamson says the following in this respect: 'But it would be foolish to respond by confining evidential probability to evidence sets which could be the total evidence possessed by a perfectly rational creature. That would largely void the notion of interest; we care about probabilities on *our* evidence' ([2002], p. 211). I am very thankful to Alan Hájek for prompting me to clarify Williamson's position in this respect.

policy will be pursued here. The discussion will assume an initial probability distribution P . P does not represent actual or hypothetical credences. Rather, P measures something like the intrinsic plausibility of hypotheses prior to investigation; this notion of intrinsic plausibility can vary in extension between contexts. ([2002], 211)

I am not committed to the view that the interpretation championed by Williamson is inadequate. For the purpose of this paper, I remain neutral on whether there is more than one adequate interpretation of evidential probabilities. However, I maintain that Williamson is wrong to argue that there is no adequate credence interpretation. While his interpretation of evidential probabilities is ingeniously and elegantly embedded in his theory of knowledge, it might very well be that a Bayesian account can coexist next to it. I am inclined to think that there is more than one adequate interpretation of evidential probabilities and that an interpretation's adequacy depends on the purpose for which the theory in question that refers to evidential probabilities is being proposed.¹² I suspect Bayesians have different purposes—and, in line with this, even different conceptions of evidential probabilities—in mind. However, to discuss this further would take us too far afield. In any case, to discuss Williamson's criticism of credence interpretations shall prove very helpful for investigations into the limits and prospects of Bayesianism.

To defend the view that there are adequate interpretations of evidential probabilities in terms of ideally rational credences, one is left with at least the following two options: first, to argue that Williamson's argument against (IAC) is problematic; second, to present an alternative interpretation and thereby show that even if Williamson's argument were not problematic, its conclusion does not establish that there is no (adequate) interpretation of evidential probabilities in terms of ideally rational credences. I shall explore both options, starting with the former.

2.3 Reviewing the Argument

Let us take a closer look at the premises (1^{IAC}) – (3^{IAC}) , which imply, together with (IAC), (\cdot^{IAC}) . (2^{IAC}) is commonly considered to be uncontroversial. According to (2^{IAC}) , if a proposition p and a proposition q are logically equivalent, then it is necessarily the case that (the credence the ideal agent s_I assigns to p , equals the credence s_I assigns to q).¹³

At first sight (1^{IAC}) and (3^{IAC}) also seem clear and uncontroversial. On a closer look, however, this is not so clear. I explain why in the following.

¹²On similar lines, Maher ([2006]) emphasizes that inductive probability plays a different role than credence. As I mentioned in Footnote 2, his conception of inductive probability is very similar to Williamson's conception of evidential probability.

¹³David Christensen ([2007]) is one of the few to think that premises such as (2^{IAC}) are to be rejected. He defends the assumption that ideal agents are rational in doubting their own logical abilities. Based on this assumption he argues that they do not assign a probability of 1 to a logical truth and, as a consequence, they do not necessarily assign the same probability to logically equivalent propositions. Here is not the place to discuss Christensen's argument. The assumption on which his argument for rejecting premises such as (2^{IAC}) is based is controversial. In the following, I assume (2^{IAC}) is unproblematic. (See (Rosenkranz [2015]) and (Smithies [2015]) for appealing ways of resisting Christensen's argument that ideal agents assign a probability less than 1 to logical truths.) After all, some epistemologists think that obeying (2^{IAC}) is essential to ideal agents.

2.3.1 ‘No One’ in (1^{IAC}) and (3^{IAC})

(1^{IAC}) and (3^{IAC}) involve b , which is supposed to be the proposition that no one has great credence in a , where a is a complex logical truth. Remember that (1^{IAC}) says the following: the evidential probability of proposition b on our evidence e^* for an agent s , is high. And (3^{IAC}) claims that it is necessarily the case that no matter what the ideal agent s_I 's total evidence is, the credence s_I assigns to the proposition $a \wedge b$ is not high. For the soundness of Williamson's argument it is crucial that the quantificatory phrase ‘no one’ is read in the same way in both premises and that ‘ b ’ likewise stands for the same proposition. Williamson certainly assumes this. However, this is not uncontroversial. The phrase ‘no one’ might well be context-dependent and ‘ b ’ express different propositions. In the present paper, I won't take a definite stand on how the phrase is to be read. It might also be that ‘no one’ is correctly understood in terms of ‘no agent’ and that the argument goes through as Williamson intended. However, in public presentations of this paper, a number of discussants have urged me to explore the possibility that the meaning of ‘no one’ is context-dependent. In the following, I argue that it is controversial what the correct reading of ‘no one’ is—and hence that Williamson's argument is in need of further motivation and is, thus, problematic.

Some have argued that the meaning of ‘no one’ is context-dependent and that (1^{IAC}) suggests a different reading of it than (3^{IAC}) : for instance, one of the premises might suggest a reading according to which it is read as ‘no human agent’, while the other might suggest a reading according to which ‘no one’ is read as ‘no human and no ideal agent’. Consequently, according to the position in question, ‘ b ’ stands for different propositions in both premises.

(1^{IAC}) and no Human Agent First, suppose that ‘no one’ is read as ‘no human agent’. Let me rephrase (1^{IAC}) accordingly. It says that the evidential probability of (no human agent has great credence in a) on our evidence e^* for an agent s , is high. e^* is our evidence, which we possess as human agents and which has some bearing on whether there are human agents that have great credence in a . So, clearly, we, human agents, can collect evidence such that the proposition that (no human agent has great credence in a) is highly probable on our evidence e^* . Therefore, given ‘no one’ is read as ‘no human agent’, (1^{IAC}) seems unproblematic.

(3^{IAC}) and no Human Agent Let us continue to suppose that ‘no one’ is read as ‘no human agent’. I rephrase (3^{IAC}) accordingly. It says that it is necessarily the case that no matter what the ideal agent s_I 's total evidence is, the credence s_I assigns to (a and *no human agent* has great credence in a) is not high. However, it is possibly the case that given our evidence e^* , the ideal agent, s_I , has great credence in the proposition that (a and *no human agent* has great credence in a). I see no reason why it is necessarily the case that the ideal agent assigns a credence in the proposition (a and *no human agent* has great credence in a) that is not high. Admittedly, it is possible that an ideal agent thinks that she is a human agent and, thus, assigns a credence in the proposition in question that is not high.¹⁴ However, since to assign such credences is not characteristic for ideal agents, this mere possibility does not suffice to establish (3^{IAC}) , which is in terms of necessity. Given ‘no one’ is read as ‘no

¹⁴Contra this, Jordan Howard Sobel ([1987]) and Peter Milne ([1991]) present Dutch book arguments in favor of ideal agents being aware that they are ideal agents. See (Christensen [2007]) and (Smithies [2015]) for criticism of this requirement.

human agent', (3^{IAC}) is—in contrast to (1^{IAC})—problematic, plausibly false, and at least in need of motivation. So what about reading 'no one' as 'no human and no ideal agent'?

(1^{IAC}), and no Human and no Ideal Agent Suppose that 'no one' is read as 'no human and no ideal agent'. Then (1^{IAC}) says that the evidential probability of (no human and no ideal agent has great credence in a) on our evidence e^* for an agent s , is high. e^* is our evidence, which we possess as human agents. While it is clear that it can have some bearing on whether human agents have great credence in a , without further argument it is not so clear that it can have some bearing on whether ideal agents have great credence in a . It is not so clear that we, human agents, might collect evidence such that the proposition that (no human and no ideal agent has great credence in a) is highly probable on our evidence e^* . The assumption that the respective evidential probability is high is in need of further motivation. Given a more objective conception of evidential probabilities it seems to come out false, and given a more subjective conception it seems to come out true. Let me explain and start with the former. It is considered to be characteristic for ideal agents that they have high credence in a . There is a conception of evidential probability that is objective in spirit and reflects this. According to it, the probability of an ideal agent having a great credence in the logical truth a is 1 independently of what our evidence might be. This is what ideal agents are *qua* being ideal. Assuming such a conception of evidential probability, the probability of the proposition (no human and no ideal agent has great credence in a) is also independent of our evidence, and is equal to 0, which makes (1^{IAC}) false. According to such an understanding, the proposition that (no human and no ideal agent has great credence in a) is not highly probable on our evidence e^* . This assumes that ideal agents exist, albeit perhaps in an abstract way. One might think that our evidence supports that there are no ideal agents, not even as abstract beings. Although I am sympathetic to this, it would be preferable to provide some motivation, which has thus far been lacking.¹⁵ There is also another conception of evidential probability, more subjective in spirit, that takes into account that from the perspective of the agent s to which the evidential probability is relativized, the proposition (no human and no ideal agent has great credence in a) is highly probable on our evidence. A possible reason behind it might be that our evidence suggests for s that no human agent has great credence in a , and that therefore one considers it also to be highly probable that no ideal agent has great credence in it, and that the conjunction (no human and no ideal agent has great credence in a) is also highly probable for s on our evidence. As a consequence the evidential probability of (no human and no ideal agent has great credence in a) on our evidence e^* for an agent s , is high. Thus, (1^{IAC}) comes out true.

In the end, whether a more objective or more subjective conception of evidential probabilities is the correct one might depend on the purpose for which the theory in question that refers to evidential probabilities is being proposed. As mentioned before, I do not intend to decide this here. I just intend to make it plausible that Williamson's argument is in need of further motivation.

(3^{IAC}), and no Human and no Ideal Agent Let us continue supposing that 'no one' is read as 'no human and no ideal agent'. Then (3^{IAC}) says that it is necessarily the case that no matter what the ideal agent s_I 's total evidence is, the credence s_I assigns to (a and

¹⁵A helpful commentary by Maria Lasonen-Aarnio made me aware that I have to make this clear.

no human and no ideal agent has great credence in a) is not high. In order for this to be necessarily so the ideal agent must be aware that she is a human or an ideal agent. It is questionable whether it is characteristic for ideal agents that they have this kind of awareness. One can plausibly save (3^{IAC}) by replacing ‘human agent’ by ‘non-ideal agent’. It seems very plausible that it is necessarily the case that no matter what the ideal agent s_I ’s total evidence is, the credence s_I assigns to the proposition (a and no non-ideal and no ideal agent has great credence in a) is not high. Presumably, the ideal agent s_I is aware of being either a non-ideal or an ideal agent and would have the self-awareness required to not assign high credence to the proposition that (a and no non-ideal and no ideal agent has great credence in a), and this is necessarily so no matter what the ideal agent s_I ’s total evidence is.¹⁶ While replacing ‘human agent’ by ‘non-ideal agent’ seems to save (3^{IAC}) , my argumentation above with respect to (1^{IAC}) and the reading of ‘no one’ as ‘no human agent’ and ‘no non-ideal agent’, respectively, would still go through and suggest that Williamson’s argument is in need of further motivation.

The previous discussion can be summarized by the following table:

Premise	No Human	No Human and No Ideal
(1^{IAC})	True	In Need of Motivation
(3^{IAC})	In Need of Motivation	True

I considered the possibility that the meaning of ‘no one’ might vary from context to context. Taking this possibility into account I rephrased ‘no one’ as ‘no human agent’ and as ‘no human and no ideal agent’, respectively. In accordance with this, I then specified the proposition b in an explicit way. Admittedly, it might turn out to be more appropriate to deal with the different contexts by restricting the domain of the quantifier from context to context differently and presenting an account of how to analyze indexicals such as ‘no one’, and then specify the relevant doxastic content accordingly.¹⁷ In the literature one finds various such accounts of quantifier domain restrictions and of analysis of indexicals. To discuss such approaches and combinations thereof would go beyond the scope of the paper here. I leave it to defendants of Williamson’s argument to present their intended accounts. My purpose here was to show that Williamson’s argument is in need of further motivation, and is thus problematic.

Even if ‘no one’ were not context-dependent, not restricted to different quantifier domains, and read in the same way in (1^{IAC}) and (3^{IAC}) , Bayesians have another strategy at their disposal to reject Williamson’s argument and endorse (IAC). I present such a strategy in the following.

2.3.2 Credence in Moore-Paradoxical Propositions and (3^{IAC})

The premises of Williamson’s argument are rendered plausible by certain assumptions about ideal agents. According to Williamson, ideal agents are not only logically and mathematically omniscient and, therefore, assign the same credence to logically equivalent propositions (see (2^{IAC})), but they also have the kind of self-awareness of their own doxastic states that prevents them from assigning high credence to a Moore-paradoxical proposition such as $a \wedge b$

¹⁶Concerning a related issue see also (Titelbaum [2013]) and Footnote 21 of this paper.

¹⁷For an instructive overview on such quantifier domain restriction, see (Stanley and Gendler Szabo [2000]).

(see (3^{IAC})). When presenting his diagnosis of his criticism of the ideal-agent-credence interpretation, he writes: ‘Presumably, a perfectly rational being must give great credence to a , be aware of doing so, and therefore give little credence to b and so to $a \wedge b$; but then its evidence about its own states would be different from ours’ ([2002], p. 210; notation adjusted).¹⁸ In the literature, there is disagreement about why it is not ideally rational to believe Moore-paradoxical propositions.¹⁹ It is not entirely clear to me which kinds of abilities and belief-forming processes ideal agents have that prevent them from having high credence in Moore-paradoxical propositions such as $a \wedge b$. However, presumably the abilities involved are different to the abilities that make them logically and mathematically omniscient. Assuming this, defenders of (IAC) could plausibly suggest we interpret the evidential probability of a proposition on an agent’s total evidence in terms of credence of an ideal agent who is logically and mathematically omniscient but lacks the kind of ability that prevents the agent from assigning high credence to a Moore-paradoxical proposition such as $a \wedge b$. Assuming such an ideal agent, one can plausibly reject (3^{IAC}) and thereby save (IAC). I consider this to be a tenable way out for advocates of (IAC). After all, ideal agents are constructed by us.²⁰ We can adjust them to our needs. However, according to this way out, the ideal agent referred to in (IAC) is only ideal with respect to the domain of logic and mathematics. Consequently, (the notion of) ideal rationality would also be restricted to the domain of logic and mathematics. Instead of stopping here and considering (IAC) to have been saved, let us look for an alternative way out that is less restrictive.

3 The Ought-Credence Interpretation

My overarching aim is to show on behalf of Bayesians how one can interpret evidential probabilities in terms of ideally rational credence and how one can maintain a tight connection between evidential probabilities and ideally rational credence even if the former cannot be interpreted in terms of the latter. Such a tight connection would be established by showing that evidential probabilities can be interpreted as ideally rational credences. In the following, I propose such a credence interpretation that does not rely on assumptions about ideal agents, and that is implicitly assumed in some discussions in the literature.

3.1 The Interpretation

I propose that instead of saving the ideal-agent-credence interpretation by theorizing about ideal agents, their credences, and the evidence they might have, one might try to avoid invoking ideal agents altogether.²¹ At the same time, one might still acknowledge that evidential

¹⁸For Williamson’s more general and different explanation of what is wrong with Moore-paradoxical conjunctions of the form ‘ a and I do not know that a ’ and ‘ a and I do not believe that a ’, see (Williamson [2002], Sect.11.3).

¹⁹See (Smithies [2016]) for a discussion of popular explanations for why it is not rational, as well as for Smithies’s own explanation. To be precise, Smithies distinguishes two kinds of Moore-paradoxical propositions. First, Moore-paradoxical propositions of the form ‘ p , but I don’t believe p ’, which is referred to as the *omissive form*, and, second, Moore-paradoxical propositions of the form ‘I believe that p , but it is not the case that p ’, which is referred to as the *commissive form*. Strictly speaking $a \wedge b$ does not correspond to either form—but certainly more to the former. However, much of what Smithies says with respect to the two kinds of Moore-paradoxical propositions in general can be applied to $a \wedge b$ as well.

²⁰That ideal agents are mere constructs is also emphasized by Michael Titelbaum ([2013], Sect. 4.2).

²¹My position with respect to theorizing about ideal agents is very much in line with Michael Titelbaum’s ([2013], Sect.4.2) insightful discussion of the use of ideal agents. Like Christensen ([2007]), Titelbaum ([2013])

probabilities are to be interpreted in terms of ideally rational credences.

A common approach to understanding ideally rational credences is in terms of credences of ideal agents. In contrast to the focus on agents that are ideal, one may focus on states that are ideal, that is, states that are how they ought to be. Thus, a second approach to understanding ideally rational credences is in terms of credences one ought to have. Up till now I have focused on the former approach. In the following, I focus on the latter, which I favour, since it renders superfluous discussions about our assumptions with respect to the meaning of ‘no one’, and the nature, existence, and abilities of ideal agents.

For the time being, I propose to interpret evidential probabilities as ideally rational credences that are understood in terms of credences one ought to have (I shall restrict this later). I think that it is safe to say that such an interpretation is already assumed by many Bayesians.²² It is certainly assumed by many Bayesians who focus on the idealness of states—as opposed to the idealness of agents.

The notion of ideal rationality is commonly considered to be normative. This is in the spirit of Bayesianism, which is also commonly considered to be a normative position. Accordingly, let me introduce what I call ‘the ought-credence interpretation’:²³

(OC) The evidential probability of a proposition p on evidence e for an agent s , $\text{Pr}_s(p|e)$, equals r if and only if it ought to be the case that the credence s assigns to p , $\text{Cr}_s(p)$, equals r if e equals s ’s total evidence, $\text{tev}(s)$.²⁴

Formally: $\text{Pr}_s(p|e) = r \leftrightarrow \mathbf{O}(e = \text{tev}(s) \rightarrow \text{Cr}_s(p) = r)$

Before I rephrase Williamson’s argument let me add some clarificatory remarks to avoid confusion. One might be tempted to object that (OC) is not adequate because it is over-

emphasizes the side-effects of the use of ideal agents and mentions Williamson’s argument against credence interpretations as an example that displays this. Titelbaum also says the following:

Williamson (2000, Section 10.1) considers what degree of belief we (given our actual evidence) should assign to the claim that ideal agents exist. On the ideal agent approach, this question is answered by considering what attitudes an ideal agent with our evidence would assign, which then becomes a question about whether the ideal agent recognizes herself as ideal and more generally what the requirements are on higher-order degrees of belief. But clearly the question of how confident we should be that ideal agents exist has little to do with higher-order strictures. ([2013], p. 74)

²²For instance, see (Bradley [2015], Ch. 5.) for a discussion of an interpretation that is in similar normative terms.

²³As a referee of this paper mentioned, it would be interesting to explore how (OC) relates to debates on the aim of minimizing expected inaccuracy, when inaccuracy is measured with the help of a strictly proper scoring rule. From the perspective of an immodest credence function, the function itself is the only credence function that one should adopt since it is the only function that minimizes expected inaccuracy. Probabilistic credence functions are immodest in the sense that they expect themselves to be the only credence function that minimizes inaccuracy.

²⁴It seems that if one takes this second approach, or way, where one theorizes about which credences one ought to have as dialectically basic and has a conception of ideal agents in mind according to which ideal agents are defined by reference to such credences, then one can keep theorizing about ideal agents. However, this way of theorizing seems very different to the one in the previous section where one takes theorizing about ideal agents as dialectically basic. The respective interpretation of evidential probability in terms of credences of ideal agents understood in this second way would have no advantage over (OC).

demanding. It equates evidential probabilities with credences one ought to have, and human agents cannot have such credences, that is, credences that obey the probability calculus. Now, roughly, the idea behind such an over-demandingness objection is based on the well-known ought-implies-can principle, which says that if one ought to do this or that, one can do this or that. Whether human agents cannot have credences that obey the probability calculus depends on the understanding of ‘can’ in question. For instance, if one understands ‘can’ in the sense of logical possibility, human agents can have credences that obey the probability calculus. It is certainly logically possible that human agents have such credences. However, if one is tempted to raise the over-demandingness objection, one might have another, more substantial understanding of ‘can’ in mind. For instance, if one understands ‘can’ in the sense of cognitively able, I think it is safe to say that most human agents cannot be logically omniscient and assign the same probability to logically equivalent propositions. Given such an understanding of ‘can’, some human agents cannot have credences that obey the probability calculus. However, even if one has a more substantial understanding of ‘can’ in mind according to which (OC) is too demanding for human agents, the over-demandingness objection poses no problem for (OC) because the ought-implies-can principle does not need to be accepted. The ought-implies-can principle is intuitively appealing but there is a reading of ‘ought’ that does not commit one to accepting it. In the following, I introduce such a reading.

Assuming an evaluative reading, or sense, of ‘ought’, the ought-implies-can principle does not hold. In the literature one distinguishes the deliberative and the evaluative readings of ‘ought’. Roughly, according to the deliberative reading, ‘ought’ means what one ought to do and in one way or another provides guidance. Schroeder specifies the evaluative reading as follows: “‘ought’ [...] has an *evaluative* sense, on which it means, roughly, that were things ideal, some proposition would be the case’ ([2011], p. 1). It is this reading I am interested in here. In the present context, I propose that ‘ought’ be read in the evaluative sense as characterized by Schroeder in the quotation above. I think the evaluative reading of ‘ought’ does not commit one to accepting the ought-implies-can principle, since it can be that if things were ideal it would be the case that *s* would do this or that although *s* cannot do this or that. This understanding of rationality in terms of ‘ought’ but without the ought-implies-can principle is in line with the following quotation by Christensen:

Clearly, we don’t want to blame anyone for failing to live up to an unattainable ideal. But there are certainly evaluative notions that are not subject to ‘ought’-implies-‘can’. I would argue that our ordinary notion of rationality is one of them: when we call a paranoid schizophrenic ‘irrational’, we in no sense imply that he has the ability to do better. ([2007], p. 5)

Given an evaluative reading of ‘ought’ we can dismiss the objection that (OC) is over-demanding.^{25,26}

²⁵In the following I shall exclusively adopt such an evaluative reading. For convenience I might sometimes even use phrases like ‘ought to do’ to express it.

²⁶For an appealing notion of ideal rationality that is considered to be evaluative in nature and is not in terms of ideal agents, see also (Titelbaum [2013], especially Chapter 4 and 5) and (Titelbaum [2015]). In a Bayesian context, Titelbaum makes explicit that he is interested in evaluations of the (evolution of) agent’s doxastic states and not evaluations of agents. Unfortunately, there is no room here to compare his notion of ideal rationality with mine. For an evaluative approach to ideal rationality that is in terms of full belief, see also (Easwaran and Fitelson [2015]).

3.2 The Rephrased Argument

Let me rephrase Williamson's argument in terms of 'ought' and by assuming (OC):

(1^{OC}) The evidential probability of proposition b on our evidence e^* for an agent s , $\text{Pr}_s(b|e^*)$, is high (that is, equal to or above the appropriate threshold t , where $t > .5$).

Formally: $\text{Pr}_s(b|e^*) \geq t$

(2^{OC}) If a proposition p and a proposition q are logically equivalent, then it ought to be the case that the credence an agent s assigns to p , $\text{Cr}_s(p)$, equals the credence s assigns to q , $\text{Cr}_s(q)$.

Formally: $p \equiv q \rightarrow \mathbf{O}(\text{Cr}_s(p) = \text{Cr}_s(q))$

(3^{OC}) It ought to be the case that, no matter what agent s 's total evidence, $\text{tev}(s)$, is, the credence s assigns to $a \wedge b$, $\text{Cr}_s(a \wedge b)$, is not high.

Formally: $\mathbf{O}(e = \text{tev}(s) \rightarrow \text{Cr}_s(a \wedge b) < t)$

(OC) and (1^{OC}) imply:

(4^{OC}) $\mathbf{O}(e^* = \text{tev}(s) \rightarrow \text{Cr}_s(b) \geq t)$

Since a is a logical truth, the following is entailed by (2^{OC}) and (4^{OC}):

(5^{OC}) $\mathbf{O}(e^* = \text{tev}(s) \rightarrow \text{Cr}_s(a \wedge b) \geq t)$

(3^{OC}) and (5^{OC}) imply:

(6^{OC}) $\mathbf{O}(e^* = \text{tev}(s) \rightarrow (\text{Cr}_s(a \wedge b) \geq t \wedge \text{Cr}_s(a \wedge b) < t))$

(6^{OC}) implies the following:

(\therefore ^{OC}) $\mathbf{O}\neg(e^* = \text{tev}(s))$

In Williamson's argument one draws the conclusion that it is not possible that the ideal agent has our evidence as total evidence, which shows that (IAC) is misguided. In contrast, in the rephrased argument here one draws the conclusion that it ought not to be the case that one has our evidence as total evidence. What this analogous rephrased argument shows is not that there is no credence interpretation (of evidential probabilities) or that there is only a misguided one, but rather that there is evidence we ought not to have (as total evidence), which, however, does not establish that (OC) is misguided. According to (OC), the evidential probability of a proposition p on evidence e for an agent s equals r if and only if it ought to be the case that the credence s assigns to p equals r if e equals s 's total evidence. It is thereby assumed that 'ought' has an evaluative reading and expresses what would be the case if things were ideal. The conclusion (\therefore ^{OC}), then, can be read as saying that it is ideal when an agent s does not have our evidence e^* as s 's total evidence. Instead of showing that the interpretation is misguided, it just makes clear that there is something wrong when we have our evidence as total evidence. It is wrong because on our evidence the evidential probability of (no one has great credence in a) is high. However, in the ideal case the evidential probability is not high since in the ideal case one has great credence in the logical truth a . Unfortunately,

things are not as they ought to be—i.e. not ideal. This, however, does not speak against the ought-credence interpretation. Embracing the conclusion $(\cdot:\text{OC})$ seems a tenable way out for Bayesians. Adding a minor revision, I refer to it as ‘the first way out for Bayesians’. The way out needs a minor revision—at least if one assumes deductive rules that conform to standard deontic logic and one understands the conditionals in (OC) as material conditionals. (OC) needs to be restricted in the following way to avoid a contradiction that results when one has evidence one ought not to have:

(OC*) If it is not the case that it ought to be the case that e does not equal s ’s total evidence, $\text{tev}(s)$, then the probability of a proposition p on some body of evidence e for an agent s , $\text{Pr}_s(p|e)$, equals r if and only if it ought to be the case that the credence s assigns to p , $\text{Cr}_s(p)$, equals r , if e equals s ’s total evidence, $\text{tev}(s)$.

Formally: $\neg\mathbf{O}\neg(e = \text{tev}(s)) \rightarrow (\text{Pr}_s(p|e) = r \leftrightarrow \mathbf{O}(e = \text{tev}(s) \rightarrow \text{Cr}_s(p) = r))$

(OC*) adds to (OC) the condition that it is not the case that it ought to be the case that e does not equal s ’s total evidence, $\text{tev}(s)$ (that is, $\neg\mathbf{O}\neg(e = \text{tev}(s))$). Without a restriction (OC) together with $(\cdot:\text{OC})$ leads to a contradiction: $(\cdot:\text{OC})$ is the same as $\mathbf{O}\neg(e = \text{tev}(s))$ and the latter entails $\mathbf{O}(e = \text{tev}(s) \rightarrow \text{Cr}_s(p) = r)$ for any p and r . This is so because $\neg(e = \text{tev}(s))$ entails $e = \text{tev}(s) \rightarrow \text{Cr}_s(p) = r$ for any p and r . $\mathbf{O}(e = \text{tev}(s) \rightarrow \text{Cr}_s(p) = r)$ together with the bi-conditional (OC) entails $\text{Pr}_s(p|e) = r$ for any p and r . The same can be shown for any probability, thus also for different r^* and r^{**} . This results in $\text{Pr}_s(p|e) = r^* \neq r^{**} = \text{Pr}_s(p|e)$ and, therefore, in a contradiction.²⁷ By restricting (OC) in the above way one avoids the contradiction. A moment’s reflection reveals that the above argumentation which assumes (OC) can be analogously applied assuming (OC*) instead. We achieve the same result, $(\cdot:\text{OC})$. As a Bayesian one can, then, endorse (OC*) and embrace that there is evidence that we ought not to have as our total evidence. A price to pay is that is not clear what to do when one has such evidence that one ought not to have.

Admittedly, (OC*) does not leave one completely satisfied when one wants more regulation for the formation of one’s credences. It is not clear what to do with evidence one ought not to have, that is, evidence one possesses when things are not ideal. However, Bayesianism as a position that merely interprets evidential probabilities in terms of ideally rational credences and still endorses all the premises (1^{OC}) – (3^{OC}) , might not be the suitable account to provide such regulation on its own; it might merely evaluate credences as ideal or not ideal. However, one might add some further principles that provide some regulation by saying what is the best thing to do when things are not ideal. In this spirit one might develop an alternative deontic logic that allows for the establishment of so-called contrary-to-duty norms. Such an alternative logic might also allow for a different understanding of the conditionals in (OC) and might allow for an alternative reading of (OC) that does not call for a revision of (OC). To discuss this here in more detail and develop such a logic is too big a task for this paper.

3.3 Reviewing the Rephrased Argument

Instead of endorsing that there is something wrong with our evidence e^* , Bayesians have the alternative of adopting a conception of evidential probabilities that makes (1^{OC}) false. In

²⁷I am grateful to a referee of this paper for the commentary that made me aware of the need to suggest a minor revision, and to Brendan Balcerak Jackson and Timothy Williamson for making a related commentary.

the following, I explore a way out for Bayesians who are not willing to accept that there is evidence one ought not to have. That said, I will adopt the original version of (OC)—no revision is required.

Before I present the second way out for Bayesians, let us take a closer look at the premises of the rephrased argument, (1^{OC}) – (3^{OC}) , and some further assumptions related to the premises. This will help us get a better grip on this way out.

3.3.1 Logical Omniscience and (2^{OC})

(2^{OC}) claims that if a proposition p and a proposition q are logically equivalent, then it ought to be the case that the credence an agent s assigns to p equals the credence s assigns to q . This premise is uncontroversial, especially given the evaluative reading of ‘ought’, according to which ‘ought’ just indicates what would be ideally the case. If one accepted (OC) in the first place, one is in some way or other committed to agreeing that an agent’s credences ought to obey the probability calculus and, therefore, committed to agreeing with (2^{OC}) . All probability functions assign the same probability to logically equivalent propositions. To reject (2^{OC}) amounts to rejecting (OC) right away. To investigate the limits and prospects of Bayesianism this does not seem an attractive option. Much the same holds for the following assumption:

(7^{OC}) It ought to be the case that the credence an agent s assigns to a , $\text{Cr}_s(a)$, is high.

Formally: $\mathbf{O}(\text{Cr}_s(a) \geq t)$

All probability functions assign probability 1 to logical truths and thus to a . In the spirit of (OC), the agent ought to assign this credence.

3.3.2 (1^{OC}) or (3^{OC}) , One Must Go

In contrast to (2^{OC}) , (3^{OC}) is not so uncontroversial. There are probability functions that assign a high probability to $a \wedge b$. (3^{OC}) excludes that they are interpreted as the credence function one ought to have. A reason for excluding such probability functions, in the context of ideal agents, is that one requires that the ideal agent has self-awareness of her own doxastic states that prevents her from having high credence in Moore-paradoxical propositions. As I mentioned before, advocates of the ideal-agent-credence interpretation might save the interpretation by interpreting evidential probabilities in terms of credences of an ideal agent who is understood as a logically and mathematically omniscient agent, who does not have the self-awareness that prevents her from assigning a high credence to the Moore-paradoxical proposition. Based on similar considerations one might reject the rephrased premise (3^{OC}) . One might argue that it is not the case that agents ought to have the self-awareness in question. However, (3^{OC}) has some intuitive appeal—certainly more so than (3^{IAC}) . From an epistemic perspective, to have this kind of self-awareness is in some way or other ideally rational. For the sake of argument, let us assume that agents ought to have the kind of self-awareness required for (3^{OC}) to hold. The idea that they ought to have this kind of awareness commits us to the following assumption, which is also what motivates (3^{OC}) —as we shall shortly have chance to observe:

(8^{OC}) It ought to be the case that (the credence an agent s assigns to the proposition a , $\text{Cr}_s(a)$, is high, if and only if the credence s assigns to the proposition b , $\text{Cr}_s(b)$, is low).

Formally: $\mathbf{O}(\text{Cr}_s(a) \geq t \rightarrow \text{Cr}_s(b) < t)$

If one already thinks that the kind of self-awareness that makes (3^{OC}) true is ideal, then it is only natural to think that it ought to be the case that if the credence s assigns to the complex logical truth a is high, then s 's total evidence contains the information that the credence s assigns to a is high. And it ought to be the case that if the total evidence contains the information that the credence the agent assigns to a is high, then the credence the agent assigns to b is low.

In the following, I show that (OC) together with (7^{OC}) and (8^{OC}) is not compatible with (1^{OC}). (7^{OC}) and (8^{OC}) imply the following:

(9^{OC}) $\mathbf{O}(\text{Cr}_s(b) < t)$

(9^{OC}) trivially yields:

(10^{OC}) $\mathbf{O}(e^* = \text{tev}(s) \rightarrow \text{Cr}_s(b) < t)$

From (10^{OC}) together with (OC) it follows that:

(11^{OC}) $\text{Pr}_s(b|e^*) < t$

(11^{OC}) is the negation of (1^{OC}). (OC), (7^{OC}) and (8^{OC}), force one to reject (1^{OC}). Therefore, if Bayesians want to save (OC), they are left with the following two options, given that (2^{OC}) is uncontroversial: the first option is to reject the assumption (8^{OC}) and the second option is to reject (1^{OC}).

Let me start with the first option. One can deny (8^{OC}). If one does so, then one is also committed to rejecting (3^{OC}) because—as mentioned—it is (8^{OC}) that supports the claim that no matter what agent s 's total evidence is, the credence s assigns to $a \wedge b$ is not high. Remember, Williamson himself writes: ‘Presumably, a perfectly rational being must give great credence to a , be aware of doing so, and therefore give little credence to b and so to $a \wedge b$ ’ ([2002], p. 210; notation adjusted). If we read the ‘must’ as a normative term, then Williamson seems to be in full agreement: (3^{OC}) is supported by (8^{OC}). From the undisputed (2^{OC}) we know that $\mathbf{O}(\text{Cr}_s(a \wedge b) = \text{Cr}_s(b))$. So if $\mathbf{O}(\text{Cr}_s(a \wedge b) = \text{Cr}_s(b) < t)$, then this must be derived from two assumptions: $\mathbf{O}(\text{Cr}_s(a) \geq t)$ and (8^{OC}). Thus, (8^{OC}) provides motivation for (3^{OC}). Assuming that the kind of self-awareness that is addressed in (3^{OC}) is ideal, rejecting (8^{OC})—and, thereby, losing this motivation for (3^{OC})—is no option. Let us focus on the second option for saving (OC).

Advocates of (OC) who want to uphold (8^{OC}) and (3^{OC}) must deny (1^{OC}), which states that the evidential probability of b on our evidence e^* for an agent s is high, and thus, given the specification at hand, that s ought to have a high credence in b on e^* . This amounts to restricting the set of evidential probability functions and the credence functions that are permitted. One might claim that (1^{OC}) is true by assumption, that any relevant interpretation of evidential probabilities should be applicable to probability functions that satisfy (1^{OC}), and

that instead of rejecting (1^{OC}) one should reject (OC). However, from a Bayesian perspective one cannot assume that all epistemically relevant interpretations of probability are applicable to probability functions that satisfy (1^{OC}). If one did so nevertheless, one would then commit oneself to a position according to which the credence in b one ought to have, or that is ideal, and the evidential probability of b on e^* —whatever the latter may be—can come apart. I am inclined to think that if one upholds (1^{OC}), one already excludes (OC). If one upholds it, this must be so because one already rejects interpretations of evidential probabilities in terms of ideally rational credences. However, in this case one might wish for an account of how evidential probabilities are related to ideally rational credences. Williamson does not focus on ideal rationality so he may not need such an account, but for those who do, such an account would be desirable.

4 Credence Gaps

As remarked at the beginning of this paper, Bayesianism—as understood here—is concerned with ideal rationality. It assumes a tight connection between evidential probabilities and ideally rational credences. This tight connection might be due to interpreting evidential probabilities in terms of ideally rational credences—as credences of an ideal agent or as credences one ought to have—or it might be due to bridge principles that connect evidential probabilities and ideally rational credences. Such a bridge principle restricts which credences are ideally rational. In this section, I focus on the latter kind of connection. In the previous sections, I focused on the former kind.

I have presented Williamson’s criticism of credence interpretations. Besides challenging Williamson’s argument itself, I also discussed two ways out for Bayesians who want to interpret evidential probabilities in terms of credences: first, one can simply embrace (\cdot^{OC}) and thereby accept that there is evidence one ought not to have. Second, one can accept that evidential probabilities and credences, respectively, are restricted in such a way that (1^{OC}) does not hold.

Bayesians who think that these ways out are too high a price for saving credence interpretations can find a third way out by replacing (OC) or (OC*) with a bridge principle that connects evidential probabilities and ideally rational credences in the following way:

(OCG) It ought to be the case that if the evidential probability of a proposition p on evidence e for an agent s , $\text{Pr}_s(p|e)$, equals r , then it is not the case that s assigns another credence to p , $\text{Cr}_s(p)$, than r if e equals s ’s total evidence, $\text{tev}(s)$.

Formally: $\mathbf{O}(\text{Pr}_s(p|e) = r \rightarrow (e = \text{tev}(s) \rightarrow \neg \exists x (\text{Cr}_s(p) = x \wedge x \neq r)))$

which is logically equivalent to:

$\mathbf{O}\neg(\exists x (\text{Cr}_s(p) = x \wedge x \neq r) \wedge \text{Pr}_s(p|\text{tev}(s)) = r)$

(OCG) is a bridge principle that expresses that it ought not to be the case that an agent s assigns a credence to a proposition p that is different to the evidential probability of p on s ’s total evidence.²⁸ This is for instance satisfied when one refrains from assigning a credence to

²⁸The scope of the ‘ought’ phrase and the ‘ \mathbf{O} ’ operator is wider in (OCG) than in (OC). While \mathbf{O} ranges over ‘ $\text{Pr}_s(p|e) = r \rightarrow (e = \text{tev}(s) \rightarrow \neg \exists x (\text{Cr}_s(p) = x \wedge x \neq r))$ ’ in (OCG), it only ranges over ‘ $e = \text{tev}(s) \rightarrow \text{Cr}_s(p) = r$ ’ in (OC).

p . By allowing for gaps of credences one can satisfy (OCG).^{29,30} In comparison to (OC) or (OC*), it is not required that, if one's total evidence equals a body of evidence one's credence in a proposition is the same as the evidential probability of the proposition given the body of evidence. Thus, by replacing (OC) or (OC*) by (OCG) one cannot derive $(\cdot:OC)$ anymore. Bayesians can agree with (1^{OC}) , and by not accepting (OC) or (OC*) but (OCG) accept that there is nothing wrong with refraining from having credences in accordance with (1^{OC}) , that is credences that correspond to the evidential probabilities in question. For this reason one cannot derive (4^{OC}) from (OCG) together with (1^{OC}) , and in a further outcome $(\cdot:OC)$ from (4^{OC}) and the other premises. By accepting (OCG) Bayesians certainly do not need to pay a high price. They can accept (1^{OC}) and at the same time not embrace $(\cdot:OC)$! The price they pay here is that they do not offer an interpretation of evidential probabilities, but merely a bridge principle that makes room for credence gaps even when evidential probabilities are defined. Accordingly, the evidential probability function is still a function, but the credences do not correspond to a probability function; rather, they correspond to a partial probability function due to the credence gaps.

5 Conclusion

Bayesianism assumes a tight connection between evidential probabilities and ideally rational credences. The tight connection might be established, first, by an interpretation of evidential probability in terms of ideally rational credences or, second, by a bridge principle that connects evidential probabilities with ideally rational credences. Such an interpretation can be understood in at least two ways. First by understanding ideally rational credences as credences of ideal agents, and second by understanding ideally rational credences as credences one ought to have. I focused first on the former understanding and reconstructed Williamson's argument (against the ideal-agent-credence interpretation, that is, (IAC)). I argued that his argument is in need of further motivation. I showed that even if it were sound it would not establish that evidential probabilities cannot be interpreted in terms of ideally rational credences. I showed this by presenting the ought-credence interpretations (OC) and (OC*). However, these interpretations come at some price: Bayesians either have to embrace that there is evidence that one ought not to have (this corresponds to the first way out for Bayesians) or embrace that evidential probabilities are restricted in such a way that does not allow that the evidential probability that no one has great credence in a is not high on our evidence, where a is a complex logical truth (this corresponds to the second way out for Bayesians). For Bayesians who do not want to pay one of the prices just mentioned I presented at the end a bridge principle that connects evidential probabilities with ideally rational credences (this corresponds to the third way out for Bayesians). This principle allows for gaps of credences where Bayesians can admit that the evidential probability that no one has great credence in a is high but there is nothing wrong with not assigning high credence in the proposition that no one has great credence in a , given one has our evidence e^* as total evidence. I have

²⁹'G' in '(OCG)' stands for 'credence gaps'.

³⁰In accordance with (OCG), presumably (2^{OC}) and (3^{OC}) should be rephrased as well. The rephrased (2^{OC}) should require that if a proposition p and a proposition q are logically equivalent, then it ought not to be the case that an agent s assigns different credences to p and to q . However, this should be satisfied when s refrains from assigning the same credence to p and to q by having gaps of credences. In the same spirit, the rephrased (3^{OC}) should be understood as being satisfied when the agent refrains from assigning a credence to the Moore-paradoxical proposition $a \wedge b$.

not argued for a specific tight connection between evidential probabilities and ideally rational credences. This is deliberate. I think that the adequacy of the particular account depends on the purpose of the specific theory of evidential probabilities and ideally rational credences which we have adopted; it thus depends on the specific Bayesian variant. I am a pluralist in this respect. An exploration of such purposes and the respective Bayesian variants must be left for a future publication.

Acknowledgments I am grateful to anonymous referees of this paper and Nathan Ballantyne, Wolfgang Barz, Peter Brössel, Kevin Dorst, Branden Fitelson, Alan Hájek, Liam Kofi Bright, Hannes Leitgeb, Maria Lasonen-Aarnio, Jim Pryor, Wolfgang Spohn, Raphael van Riel, Ben Young, and Timothy Williamson for very constructive commentaries on (parts of) previous versions of this paper. I am also grateful to David Christensen, Philip Ebert, Kate Elgin, Insa Lawler, Ram Neta, Declan Smithies and Mike Titelbaum for insightful and helpful commentaries on a related paper. Their commentaries were also helpful for this paper. I have presented previous versions of the present paper at various events, where I received helpful criticism from Brendan Balcerak Jackson, Magdalena Balcerak Jackson, Julien Dutant, Kenny Easwaran, Clayton Littlejohn, Richard Pettigrew, and Luis Rosa. I am also grateful to the participants of the following workshops and conferences: *Statistical Evidence in Epistemology and the Law* in Glasgow; the *Research Colloquium of the Chair for Theoretical Philosophy* at the University of Duisburg-Essen; the *British Society for the Philosophy of Science Annual Conference* in Manchester; the *Ninth European Congress of Analytic Philosophy* in Munich; the *Invited Symposium: Ideal and Non-ideal in Epistemology and Political Theory* at the APA Pacific Division 2018; the workshop *Knowledge, Belief, and Probability* in Hamburg, and the *Epistemology Seminar 2018*.

Funding for this research was provided by the Austrian Science Fund (Nr. J 3783 - G24).

References

- [1] Bradley, D. [2015]: *A Critical Introduction to Formal Epistemology*, Bloomsbury.
- [2] Brössel, P. [2012]: *Rethinking Bayesian Confirmation Theory*, PhD Thesis, University of Konstanz.
- [3] Brössel, P. and Eder, AM. A. [2014]: ‘How to Resolve Doxastic Disagreement’, *Synthese*, **191**, pp. 2359–81.
- [4] Chisholm, R. [1963]: ‘Contrary-to-Duty Imperatives and Deontic Logic’, *Analysis* **24**, pp. 33–6.
- [5] Christensen, D. [2007]: ‘Does Murphy’s Law Apply in Epistemology? Self-Doubt and Rational Ideals’, *Oxford Studies in Epistemology*, **2**, pp. 3–31.
- [6] Easwaran, K. and Fitelson, B. [2015]: ‘Accuracy, Coherence, and Evidence’, *Oxford Studies in Epistemology* **5**, pp. 61–96.
- [7] Eder, AM. A. and Brössel, P. [forthcoming]: ‘Evidence of Evidence as Higher Order Evidence’, in M. Skipper and A. Steglich-Petersen (eds), *Higher-Order Evidence: New Essays*, Oxford University Press.
- [8] Hájek, A. [2012]: ‘Interpretations of Probability’, in E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, available at <<https://plato.stanford.edu/archives/win2012/entries/probability-interpret/>>.

- [9] Levi, I. [1980]: *The Enterprise of Knowledge*. MIT Press.
- [10] Lewis, D. [1973]: *Counterfactuals*, Harvard University Press.
- [11] Maher, P. [2006]: ‘The Concept of Inductive Probability’, *Erkenntnis*, **65**, pp. 185–206.
- [12] Milne, P. [1991]: ‘A Dilemma for Subjective Bayesians – and How to Resolve It’, *Philosophical Studies*, **62**, pp. 307–14.
- [13] Rosenkranz, S. [2015]: ‘Fallibility and Trust’, *Noûs*, **49**, pp. 616–41.
- [14] Schroeder, M. [2011]: ‘Ought, Agents, and Actions’, *Philosophical Review*, **120**, pp. 1–41.
- [15] Sobel, J.H. [1987]: ‘Self-Doubt and Dutch Strategies’, *Australasian Journal of Philosophy*, **65**, pp. 56–81.
- [16] Smithies, D. [2015]: ‘Ideal Rationality and Logical Omniscience’, *Synthese*, **192**, pp. 2769–793.
- [17] Smithies, D. [2016]: ‘Belief and Self-Knowledge: Lessons from Moore’s Paradox’, *Philosophical Issue*, **26**, pp. 393–491.
- [18] Stanley, J. and Gendler Szabo, Z. [2000]: ‘On Quantifier Domain Restriction’, *Mind & Language*, **15**, pp. 219–61.
- [19] Stalnaker, R. [1968]: ‘A Theory of Conditionals’, in: N. Rescher (*ed*), *Studies in Logical Theory*. Basil Blackwell Publishers, pp. 98–112.
- [20] Titelbaum, M. G. [2013]: *Quitting Certainties. A Bayesian Framework Modeling Degrees of Belief*, Oxford University Press.
- [21] Titelbaum, M. [2015]: ‘Rationality’s Fixed Point (or: in Defense of Right Reason)’, *Oxford Studies in Epistemology*, **5**, Oxford University Press, pp. 253–94.
- [22] Williamson, T. [2000]: *Knowledge and its Limits*, Oxford University Press.
- [23] Williamson, T. [2002]: *Knowledge and its Limits*, Oxford University Press. [Paperback]
- [24] Yap, A. [2014]: ‘Idealization, Epistemic Logic, and Epistemology’, *Synthese*, **191**, pp. 3351–66.