

The Nature of Awareness Growth

Chloé de Canson

Abstract. Awareness growth—coming to entertain propositions of which one was previously unaware—is a crucial aspect of epistemic thriving. And yet, it is widely believed that orthodox Bayesianism cannot accommodate this phenomenon, since that would require employing supposedly defective catch-all propositions. Orthodox Bayesianism, it is concluded, must be amended. In this paper, I show that this argument fails, and that, on the contrary, the orthodox version of Bayesianism is particularly well-suited to accommodate awareness growth. For it entails what I call the *refinement view*, which allows us to capture that awareness growth consists in the increase of one’s capacity of discernment.

0.

Subjects often undergo what is known as *awareness growth*:¹ they come to entertain propositions of which they were previously unaware. Examples of this phenomenon abound: scientists formulate new theories, students discover new ways the world might be, people remark that they “didn’t know that was an option”. Such an epistemic shift does not seem irrational; instead it is clearly a crucial aspect of epistemic thriving, and may even constitute a form of learning. This suggests that an adequate account of rationality must feature awareness growth. But orthodox Bayesianism is widely taken to be unable to accommodate this phenomenon, because doing so would require the inclusion of catch-all propositions, which are thought to be

Published in the *Philosophical Review* 123(1): 1–32, 2024. This is an uncorrected manuscript: please cite the published version, available [here](#).

¹ In Bayesian philosophy of science, the phenomenon is sometimes referred to as that of *new theories* after Glymour (1980). I use the term from formal epistemology and decision theory because it is more general: it encompasses not only the formulation of new scientific hypotheses, but also more mundane phenomena which are structurally identical.

defective. As a result, many have concluded that the orthodox version of Bayesian epistemology must be rejected. This is the *argument from awareness growth*.

My aim in this paper is to show that the argument from awareness growth is unsuccessful in undermining orthodoxy, and that orthodoxy is actually very well-suited to accommodate awareness growth. So, to the extent that Bayesians want to theorise awareness growth, they would do well to adopt an orthodox position. I proceed as follows. I begin by reconstructing the argument from awareness growth (§1). Then, I show that its core premise—that endeavours involving catch-all propositions must be rejected—is unfounded (§2). In fact, I argue, the whole argument relies on an erroneous conception of the relationship between orthodoxy, awareness growth, and catch-all propositions (§3). It follows that the argument against orthodoxy from awareness growth fails. In fact, I argue, orthodoxy can actually accommodate awareness growth in an elegant and insightful way, with what I call the *refinement view*, on which awareness growth consists in the increase of one’s capacity of discernment (§4). I conclude that, if anything, the requirement to accommodate awareness growth counts in favour of, and not against orthodoxy (§5).

1.

Bayesians are characterised by a commitment to the claim that an agent’s epistemic state may be represented by a function $p : \mathcal{A} \rightarrow [0, 1]$, which assigns a degree of belief, or *credence*, to each proposition that the agent entertains. Thus an agent’s epistemic state can be thought of as comprising two components: what we might call the *awareness* component—the propositions of which the agent is aware, represented by the set \mathcal{A} ; and what we might call the *credal* component—the degree of confidence that the agent has in each of the propositions $A_i \in \mathcal{A}$, represented by the shape of the function p .

Defined as such, Bayesianism is a broad church. Throughout this paper, I shall be concerned with the orthodox version thereof (“*orthodoxy*”), which I shall define in terms of two characteristic features: a *regulative* one about the norms that govern the credal component of epistemic states, and a *hermeneutic* one about the interpretation of the Bayesian formalism. Let me begin with the former. Credence-governing norms (or, norms of rationality) have been the central concern of most Bayesian epistemology for the past hundred years; most prominently, Bayesians have debated how confident agents ought to be in various propositions, and how agents ought to revise their degrees of confidence in various propositions upon acquiring new evidence. For a Bayesian view to count as orthodox, it must maintain that (at least) two norms govern an agent’s epistemic state. The first, *probabilism*, asserts that the function p

representing a rational agent’s epistemic state must be a probability function; and the second, *conditionalisation*, asserts that upon learning a proposition $E \in \mathcal{A}$, the agent must update her credence function $p(\cdot)$ to a new credence function $p'(\cdot) = p(\cdot|E)$. The wealth of arguments for these norms has no doubt played a role in cementing them as part of the orthodoxy.²

The second (hermeneutic) feature of orthodoxy concerns the interpretation given to the formalism; in particular, to two features thereof. Probabilism entails what we shall call the *Boolean norm*, according to which the set \mathcal{A} that represents the awareness component of the agent’s epistemic state must form a Boolean algebra.³ This norm entails, among other things, that the algebra must contain the propositions we might call *trivial*, represented as Ω and \emptyset ; and propositions we might call *non-trivial*, represented as A_1, A_2, \dots . Probabilism also entails that the credal component of the agent’s epistemic state must be represented by a credence function $p : \mathcal{A} \rightarrow [0, 1]$, which assigns to each proposition a credence between 0 and 1. Let us mark the difference between *extremal* credences (credence 1 and credence 0), and *non-extremal* credences (other credal values between 0 and 1).

For a version of Bayesianism to count as orthodox, it must provide a particular interpretation of the trivial/non-trivial distinction, and of the extremal/non-extremal distinction: the latter distinction must be interpreted as the distinction between absolute certainty and less-than-absolute-certainty; and the former distinction must be given a *modal interpretation*. Let me explain. In the orthodox literature, the two most widely discussed contenders for the interpretation of the trivial/non-trivial distinction hold that it marks (i) the distinction between logically necessary and logically contingent propositions; or (ii) the distinction between metaphysically necessary and metaphysically contingent propositions (Hájek, 2012, ms; Easwaran, 2014; Mahtani, forthcoming) . I propose elsewhere a third option, namely that it marks (iii) what I call the *poric*⁴ distinction between the determinable *a priori* and the determinable *a*

² The two most prominent types of arguments for these norms are so-called *Dutch Book arguments* (such arguments for probabilism originate with Ramsey (1926) and de Finetti (1937); those for conditionalisation with Teller (1973) and Lewis (1999)) and *accuracy arguments* (see Joyce (1998) and Leitgeb and Pettigrew (2010a,b) on probabilism; and Greaves and Wallace (2006) and Briggs and Pettigrew (2020) on conditionalisation).

³ Orthodoxy entails the Boolean norm because it is partly constituted by a commitment to probabilism—the claim that agents’ epistemic states are representable by a probability function—and it is part of the definition of a probability function that it is defined on a Boolean algebra.

⁴ The term, after the Greek $\pi\acute{o}\rho\omicron\varsigma$ which refers to the means to one’s ends, points to the means one must employ (pure reason, the senses) to determine the truth-value of a proposition.

posteriori (Canson, ms).⁵ While there are differences between these interpretations which matter for various reasons and purposes, what will matter for the project of this paper is what they have in common: they all take the trivial/non-trivial distinction to mark differences in (logical/metaphysical/poric) *modal* status. So, I shall use the terms *necessary* and *contingent* as an umbrella term to encapsulate the three.⁶

So, a version of Bayesianism shall be orthodox to the extent that it maintains that probabilism and conditionalisation are norms of rationality (the regulative feature), and that the extremal/non-extremal and trivial/non-trivial distinctions are to be interpreted as the certainty/uncertainty and necessary/contingent distinctions (the hermeneutic feature). Note that, defined as such, orthodox Bayesianism is not maximally specific: there are multiple particular views that count as orthodox—views that may differ over, for instance, whether there are norms of rationality beyond probabilism and conditionalisation, or over how to interpret the rest of the formalism.

Having characterised orthodoxy, let us turn to awareness growth. By contrast with the credal component of epistemic states, there has been very little discussion of the awareness component in philosophy. Instead, the almost universal assumption has been that the set of propositions which the agent entertains (that is, the set of propositions in which the agent has an epistemic attitude) is fixed and exogenously supplied. This paper renounces this assumption following an old tradition (Shimony, 1970; Salmon, 1990; Earman, 1992), dormant and recently revived (Wenmackers and Romeijn, 2016; Bradley, 2017; Mahtani, 2021a; Steele and Stefánsson, 2021a,b).⁷ I start, as they do, from the fact that awareness growth occurs,

⁵ For a discussion of other options, see Tang (2012). For a discussion of how the variety of these interpretations relates to the arguments for probabilism, see Mahtani (2021b).

⁶ Many critics of orthodoxy have rejected it for its adoption of the modal interpretation of the trivial/non-trivial distinction. For instance, Hacking (1967) was an early critic of the logical interpretation, for entailing that agents ought to have credence 1 in all logical truths. Chalmers (2011) and more recently Mahtani (forthcoming), have criticised the metaphysical interpretation for entailing that agents ought to have credence 1 in all metaphysically necessary propositions, including those that are determinable only a posteriori—propositions such as “Hesperus is Phosphorus”—see fn. 25. This line of criticism is not directly related to the target of this paper, since the question with which I am concerned here is that of whether awareness growth poses a distinctive challenge to Bayesian orthodoxy. So, those who already reject orthodoxy for other reasons—perhaps because they reject the modal interpretation of the trivial/non-trivial distinction—will not be immediately affected by my intervention.

⁷ See also Henderson et al. (2010), Hill (2010), Carr (2015), Pettigrew (forthcoming), and Roussos (ms).

and that it occurs routinely. As Earman puts it, “new observations, even of familiar scenes; conversations with friends; idle speculations; dreams—all of these and more are constantly introducing heretofore unarticulated possibilities” (1992, p. 198). It follows that, if Bayesianism is to be an adequate theory of epistemic rationality, it must *accommodate* awareness growth; that is, it must contain descriptions of the phenomenon, so that agents who undergo this type of epistemic shift can be adequately represented.

An adequate accommodation of awareness growth will also serve to properly articulate a question on which more ink has been spent: that of how to *regulate* awareness growth; that is, of which norms of rationality there are, if any, for epistemic changes involving awareness growth.⁸ Several interventions surrounding these regulative norms have, unsurprisingly, led to deliberations on how the phenomena should be represented in the first place.⁹ Indeed, since the very formulation of the regulation question requires an answer to the accommodation question, we must have an adequate representation of these shifts. But it is widely thought that orthodox versions of Bayesianism cannot do this; that there is no satisfactory way to accommodate awareness growth therein.

To begin to see what the supposed problem is, it will be useful to distinguish three types of awareness growth as thought of pre-theoretically.¹⁰ Consider the following *book case*, which I will use as a running illustration throughout the rest of this paper. An agent has been gifted an as-yet-unwrapped book, of which they ponder the genre: they actively entertain that the book might be a novel or a collection of poems. Furthermore, they have, let us assume, never heard of short stories; we can call their state of awareness “*limited*” following Steele and Stefánsson (2021a). This agent is aware of the propositions we might denote as *novel* and *poems*, but unaware of *stories*. There are three types of awareness growth that they might undergo. Firstly, they might come to entertain that the book might be a paperback or a hardback. They would then entertain four possibilities: paperback novel, hardback novel, paperback poetry, and

⁸ This discussion has mostly taken place around the norm labelled *Reverse Bayesianism* by Karni and Vierø (2013, 2015) according to which agents ought not modify the ratios of their credences in propositions they already entertained. The norm is defended by e.g. Wenmackers and Romeijn (2016), Bradley (2017), and Roussos (msa), and rejected by Carr (2015), Mahtani (2021), Steele and Stefánsson (2021a,b), and Pettigrew (forthcoming).

⁹ Mahtani (2021), Pettigrew (forthcoming) and Roussos (msa) all argue that whether Reverse Bayesianism holds depends sensitively on the way awareness growth is represented; though they draw different conclusions about what that entails for Reverse Bayesianism.

¹⁰ I say “pre-theoretically” because, although these distinctions might be appealing at this early stage, and are implied in much of the literature, we will see in §4 that they do not withstand theorising.

hardback poetry. I call this *orthogonal* awareness growth: the newly entertained propositions are logically independent from the old ones. Secondly, they might come to consider that, if the book is a book of poems, it might be a book of prose poems or a book of poems in verse. They would then entertain three possibilities: novel, prose poems, and verse poems. I call this *internal* awareness growth: the newly entertained propositions jointly entail a single old one. Thirdly and finally, the agent might come to consider that the book also might be a collection of short stories. They would then entertain three possibilities: novel, poems, and stories. I call this *lateral* awareness growth: the newly entertained proposition is inconsistent with the old actively entertained ones.

How can the orthodox Bayesian model these epistemic events? As mentioned above, the orthodox Bayesian is committed to the Boolean norm, and therefore represents the propositions in which an agent has epistemic attitudes by a Boolean algebra \mathcal{A} . A growth in awareness must be represented by an enlargement of this algebra to another \mathcal{A}' . And a particularly perspicuous way to enlarge an algebra is by *refinement*: in such cases, \mathcal{A}' is a fine-graining of \mathcal{A} ; that is, for all $A_i \in \mathcal{A}$, $A_i \in \mathcal{A}'$. In more informal terms, refinement consists in the splitting of one or more propositions. This operation is illustrated in the figure below.




FIGURE HERE, CAPTION BELOW

Refinement of an algebra of propositions

As far as I know, there is a consensus in the literature that orthogonal and internal growths of awareness should be modelled by refinement. So, for instance, the agent who comes to entertain that the poems may be in verse or prose should be modelled such that the element *poems* of their old algebra \mathcal{A} corresponds to the union of the elements *verse* and *prose* of their new algebra \mathcal{A}' . What about lateral growths of awareness?

This is where the problem supposedly lies for orthodoxy. It has long been recognised that if it is to accommodate lateral awareness growth, orthodoxy must model it by refinement, and this requires the inclusion of what Shimony (1970) has called a *catch-all proposition* in the

agent’s algebra.¹¹ Roughly, a catch-all proposition is a proposition which expresses something like “a possibility beyond those I am entertaining”; it is the complement of the union of all propositions we can intuitively identify as *actively entertained*. Let me explain why orthodoxy requires such propositions to model lateral awareness growth. A constitutive feature of orthodoxy is the commitment to what I have called the *modal* interpretation of the Bayesian formalism, on which the sample space Ω represents the (logically, metaphysically, or porically) necessary proposition. And what is (logically, metaphysically, or porically) necessary does not change. So, on orthodoxy, the sample space Ω must remain fixed, such that the modelling of awareness growth cannot be done by what has been called “expansion” (Steele and Stefánsson, 2021a,b) of the sample space—where a sample space is replaced with another, larger one. And if the sample space must remain fixed as the agent’s awareness grows, there must have been a portion of the initial sample space corresponding to something other than the propositions the agent then actively entertained. This is the catch-all proposition—the complement of the actively entertained propositions against the modal domain of discourse. And, everyone agrees: there is an important concern with catch-all propositions.

The putative problem is that, unlike other propositions, it is uniquely difficult to rationally assign a credence to the catch-all proposition. Salmon calls this “the utter intractability of the [probability of] the catchall” (1990, p. 191). The argument for this claim is what I shall call the *defectiveness argument*, and it appears widely in the literature. Bradley asks rhetorically: “given that we don’t know anything about the prospects that we are potentially unaware of, on what basis are we to determine . . . what probability we should assign to the catch-all prospect?” (2017, p. 255). Wenmackers and Romeijn state that “since the catch-all is not based on a scientific theory, the usual . . . considerations . . . for assigning it a prior . . . do not come into play here” (2016, p. 1234). According to Henderson et al., the modelling of lateral awareness growth by refinement “is an unsatisfactory solution since there is no particularly principled way to decide how much initial probability should be assigned to the catchall” (2010, p. 190). The consensus is such: catch-all propositions are required for orthodoxy to accommodate lateral awareness growth, and they are defective.

Commentators have all reacted in the same way—by rejecting orthodoxy. The most popular way to reject orthodoxy has been to reject its hermeneutic feature by relativising the space of possibilities to the propositions that the agent actively entertains, in such a way that catch-all propositions are not included in the set of propositions in which agents have a cre-

¹¹ For work following on Shimony, see Glymour (1980), Salmon (1990), Earman (1992), Maher (1995), Zynda (1995).

dence. Shimony adopts this strategy, and claims that “a methodological principle of great importance, which permits probability to be a manageable instrument in scientific inference, ... is that *the individual investigation delimits an area in which probabilities are calculated*. ... The conditions of a single investigation establish a kind of ‘local’ universe of discourse within which calculations strictly governed by the axioms of probability can be performed” (1970, p. 99). In a similar vein, Salmon writes that “since the catchall is not a bona fide hypothesis, ... we need not try to calculate its ... probability” (1990, p. 192). Bradley concludes from his discussion of catch-all propositions that “only a slight tweak is required in the representation of a Bayesian-rational agent to allow for less than full awareness, in the form of a relativisation of her attitudes to her subjective domain of awareness” (2017, p. 256).¹² Steele and Stefánsson develop the idea of a space of possibilities relativised to what they call an “awareness context” in great detail (2021a, 2021b).¹³ And so on. Others have reacted instead by rejecting the regulative feature of orthodoxy. For example, Wenmackers and Romeijn (2016) suggest rejecting probabilism: they argue that the catch-all proposition should be assigned either an

¹² Bradley’s model actually includes two layers: he invokes “a background Boolean algebra of propositions, which we can think of as the set of all possibilities, or the set of all possibilities that the modeller can conceive of”, which he contrasts with a smaller Boolean algebra giving “the set of possibilities of which the agent is aware” (2017, p. 256). Thus, although he appeals to an algebra that looks like the one posited by orthodoxy, he is clear that the algebra which matters for the purposes of representing limited awareness and awareness growth—the algebra on which the agent’s attitudes are defined—is, on his model, relativised to the agent’s awareness state.

¹³ Steele and Stefánsson write that “an agent’s awareness context is defined by a set ... of basic propositions *of which she is aware*” (2021a, p. 1223, emphasis added), from which one can “generate a Boolean algebra ... in the usual way” (2021a, p. 1224; see also 2021b, p. 33). Thus, the agent’s algebra is relativised to their awareness state, such that “when the agent’s set of basic propositions changes upon awareness growth, so too does her possibility space” (2021a, p. 1214). In their book, they do consider employing a “subjective catch-all” to represent those agents who “anticipate” awareness growth (2021b, p. 36). But, they insist, “these types of catch-all are not meant to be interpreted as a metaphysically or logically universal complement to whatever the agent is aware of” (ibid): they are not, that is, the kind of catch-all that would have to be employed on a non-relativised, orthodox version of Bayesianism. Moreover, they are restricted to a special class of agents with limited awareness, those who anticipate awareness growth, whereas orthodoxy requires the modelling of all such agents with a catch-all propositions.

“indefinite” credence or no credence at all.¹⁴¹⁵ The merits of these moves shall not concern us here, though there is much to be said about them. Rather, what should transpire from this discussion is just how serious these authors take the problem with awareness growth to be. For them, it is a reason to reject orthodoxy.

Let us call this argument against orthodoxy the *argument from awareness growth*. Its structure is as follows:

1. Orthodoxy must accommodate lateral awareness growth.
2. This requires the employment of catch-all propositions.
3. Endeavours that require catch-all propositions must be rejected.
4. So, orthodoxy must be rejected.

In the next two sections, I examine and refute this argument. In §2, I examine the defectiveness argument, sketched above, which is supposed to establish premise 3. I argue that it fails: there is no reason to reject endeavours that require the employment of catch-all propositions. And in §3, I study the general structure of the argument from awareness growth. I show that it encodes a particular view about the relationship between orthodoxy, awareness growth, and catch-all propositions, which is false. This not only constitutes an additional reason to reject the argument, but it also paves the way for understanding the proper conceptual relations between orthodoxy, awareness growth, and catch-all propositions, and so, in turn, for seeing how to accommodate awareness growth on orthodoxy. This is what I do in §4.

¹⁴ Their former proposal to assign an “indefinite” credence to the catch-all proposition suggests another possible approach: one might assign an *imprecise probability* to the catch-all proposition. Since my aim in this paper is to show that the features of the catch-all proposition do not mandate a departure from orthodoxy provided one is already committed to it, I will not explore this option here. But it certainly merits attention; for an overview of the literature on imprecise probability, see Bradley (2019). See also fn. 17.

¹⁵ Their latter proposal to assign no credence to the catch-all proposition is strikingly close to the other type of orthodoxy rejection: the relativisation of the domain of discourse to the actively entertained propositions. Their proposal is essentially to “condition[al]ise] all evaluations on” the disjunction of all actively entertained propositions, such that we can “express the agent’s opinions concerning the relative probability of [any two actively entertained propositions] without saying anything . . . about the absolute probability that they have.” In effect, they claim, this amounts to proceeding “relative to a theoretical context” (Wenmackers and Romeijn, 2016, pp. 1241–2).

2.

Everywhere that catch-all propositions are mentioned, the argument is made that they are defective. This *defectiveness argument* is supposed to establish premise 3 of the argument from awareness growth: that endeavours involving catch-all propositions must be rejected. In this section, I expose, examine, and refute this argument.

The supposed defect with catch-all propositions is that “there is no particularly principled way to decide how much . . . probability should be assigned to the[m]” (Henderson et al., 2012, p. 190). Indeed, Bradley asks, “given that we don’t know anything about the [propositions] that we are potentially unaware of, on what basis are we to determine what . . . probability we should assign to the catch-all [proposition]?” (2017, p. 255). The point is also made by Wenmackers and Romeijn:

Whereas the other [propositions] come with a—possibly very intricate—theoretical background story, the catch-all has no content other than “none of the explicitly formulated hypotheses”. So, [it] is the [complement of actively entertained propositions] and that is all that can be said about it. [It follows that] it is not sensible to assign any definite [probability to it.] Since the catch-all is not based on a scientific theory, the usual . . . considerations . . . for assigning it a [probability] do not come into play here. (2016, pp. 1232-4)

This line of thought can be traced back to Salmon:

The point to be emphasised . . . is the utter intractability of the [probability of] the catchall. The reason for this difficulty is easy to see. Whereas the seriously considered candidates are bona fide hypotheses, the catchall is a hypothesis only in a Pickwickian sense. It refers to all of the hypotheses we are *not* taking seriously, including all those that have not been thought of as yet; indeed, the catchall is logically equivalent to their disjunction. . . . To try to evaluate the [probability of the catch-all proposition] is something we cannot do with any reliability. (1990, p. 191)

The defectiveness argument is beginning to come into view. Catch-all propositions are what we shall call *unintelligible*, such that we “don’t know anything about them” (Bradley), “all that can be said about” them is that they are the complement of actively entertained propositions (Wenmackers and Romeijn), they are not “bona fide” propositions, but propositions “only in the Pickwickian sense” (Salmon). Or, in the words of Steele and Stefánsson, “the agent has

no idea how to specify the[ir] content” (2021a, p. 1226), and “cannot (yet) articulate” them (2021b, p. 34). It follows from this that they are, to use Salmon’s word, “*intractable*”: because of their unintelligibility, “there is no particularly principled way” to assign a probability to them (Henderson et al), there is no “basis” on which to “determine what probability to assign” to them (Bradley), “it is not sensible to assign any definite” probability to them (Wenmackers and Romeijn), “to evaluate” their probability “is something we cannot do with any reliability” (Salmon).

But this link between unintelligibility and intractability, which constitutes the heart of the defectiveness argument, is rather underdeveloped in commentaries on catch-all propositions—presumably because it is taken to be obvious. Sometimes, it is drawn by mere rhetorical appeal: *how*, asks Bradley in the citation above, could we *possibly* determine how confident to be in a catch-all proposition, if we don’t know what it says? Other times, its exposition takes the shape of an illustration from the history of science (Shimony, 1970; Glymour, 1980; Salmon, 1990; Earman, 1992). For instance, consider physicists at the end of the nineteenth century, entertaining a pair of classical theories. They might have preferred one over the other if it was a better fit with the available evidence. But what reason could they have had for thinking that another, heretofore unformulated theory was true instead?

The illustration is suggestive. The scientists therein prefer one theory over the other because it fits better with the evidence. So, it is their respective fit with the evidence that makes the actively entertained propositions tractable. This intimates an elucidation of the link between a catch-all proposition’s unintelligibility and its intractability. Since the catch-all proposition is unintelligible—since we do not know what it expresses—we cannot determine its relation to the evidence. And since the assignment of a probability to a proposition is supposed to be done on the basis of its relation to the evidence, it follows that we cannot assign it a probability: it is intractable. Interestingly, Salmon as well as Wenmackers and Romeijn insist on the difficulty of determining the *likelihood* of the agent’s evidence on the catch-all proposition—the probability of the evidence the agent has given that the catch-all proposition is true. This suggests that this is how they too conceive of the link.¹⁶

There remains one last aspect of the defectiveness argument to comment on. The scientists in the illustration above have no trouble determining how confident to be in one actively entertained theory relative to another. The problem lies only with determining the probability of the catch-all proposition. This makes sense, of course, if it is unintelligibility that leads to

¹⁶ Thanks to Hylke Jellema and Jan-Willem Romeijn for very helpful conversations that helped me put the point in these terms.

intractability. Thus Wenmackers and Romeijn claim that “the ratio of two . . . probabilities assigned to explicit hypotheses can still be obtained”, even though the probability of the catch-all proposition cannot (2016, p. 1239). This point plays a crucial role in the argument as a whole: the defectiveness argument is supposed to establish that endeavours must be rejected *to the extent* that they involve catch-all propositions; they must be rejected when *and only when* they involve them. So whatever deficiency there is must be specific to and general for catch-all propositions; in other words, it must be *distinctive* of them. This is sometimes captured by the assertion that what is really intractable is the *ratio* of the catch-all proposition’s probability to the probability of its negation. On the picture painted by the defectiveness argument, it is categorically harder to determine how confident one ought to be in the catch-all proposition relative to its negation, than it is to determine how confident one ought to be in one actively entertained proposition relative to another.¹⁷

The structure of the defectiveness argument is now fully apparent. The characteristic feature of catch-all propositions is their unintelligibility. This feature, because it makes it impossible to determine how they relate to the evidence, makes them intractable. And their intractability constitutes a reason not to employ them. In sum, the defectiveness argument can be stated thus:

- i. Catch-all propositions are unintelligible.
- ii. Unintelligible propositions are intractable.
- iii. Endeavours involving intractable propositions must be rejected.
- 3. So, endeavours that require catch-all propositions must be rejected.

The crux of the argument is clearly premise **ii**, so this is where I shall start.

The premise is defended, as we saw, with examples from the history of science in which it is intuitively difficult to assign a probability to the catch-all proposition. Our physicists prefer one actively entertained theory over another because of its fit with the evidence, but have no idea how to evaluate the possibility that an altogether different theory might be true. However, examples to the opposite effect can easily be constructed. For instance, suppose that our book-receiver knows their book-gifter to be an experimental literature enthusiast. They might then have a good reason to believe that the book is likely not to be as banal as a mere novel or poetry collection; that it is much more likely indeed to be *something else* instead. Or suppose that our physicists are familiar with what Poincaré describes as the long history of scientific theories “abandoned after a few years of favor, the wreckage of one theory

¹⁷ Thanks to Franz Dietrich and to an anonymous referee for helping me see the significance of this ratio in the typical understanding of the problem.

piled on the rubble of another”. They might then be perfectly justified in “anticipating that today’s fashionable theories will soon give way, and from this conclude that these theories are completely futile” (1905/2018, p. 115).¹⁸ These counterexamples show that the intuitions which supposedly make premise **ii** so appealing are in fact much less stable than is suggested by the brevity—sometimes even rhetoricality—of their exposition.

And more importantly, they point to a diagnosis of where the argument for the premise that I outlined above goes wrong. This argument moves from the postulate that catch-all propositions are unintelligible to the claim that we “don’t know anything about the[m]” (Bradley 2017, p. 255), including about their relation to the evidence—this is why they are intractable. But a proposition’s being unintelligible simply does not entail that we can’t know or say anything informative about it. Indeed, as the two counterexamples show, there is plenty that we can know about catch-all propositions, even if we do not know their content. And in some cases, what we can know about them is precisely what we need to know in order to determine how they relate to our evidence. Not knowing what a non-novel non-poetry book might be is precisely what allows our book-receiver to infer that their friend who is well-versed in experimental literature is likely to have gifted them such a book. A theory’s not yet having been formulated is precisely what undergirds the pessimistically-minded physicist’s confidence in it. So, because a proposition’s being unintelligible does not entail knowing nothing of it, nothing systematic about tractability follows from unintelligibility. Premise **ii** is false, and the defectiveness argument fails.

But the failure of premise **ii** does not exhaust the problems with the defectiveness argument. For the argument requires, as I put it at the start of the section, that there be a characteristic feature of catch-all propositions in virtue of which they can be defective: what I have called their *unintelligibility*. This notion, when applied to the kind of example we have been considering, seems clear and intuitively appealing. The book-receiver knows how to specify the content of *novel*, but not of *other*; the physicist finds the formulated theories intelligible in a way that unformulated ones are not (indeed this feels almost tautological).

¹⁸ The argument contained in this famous passage has come to be known as the *pessimistic meta-induction*—for the canonical discussion, see Laudan (1981). On the basis of the now-recognised falsity of previously successful theories, this argument seeks to establish that we should have a very low confidence in our currently formulated theories, relative to the theories that have yet to be articulated. So, if successful, it directly challenges the defectiveness argument applied to the kind of example of scientific theories at hand here. Indeed, the pessimistic meta-induction’s conclusion directly entails that it is easy to determine how confident physicists ought to be in the negation of the theories they entertain: in a word, *very*.

But intelligibility is recalcitrant to precise theoretical treatment. For how might it be precisely stated? One option is that the agent be able to enumerate all the possible instances of a particular proposition. But this is clearly too strong: our book-receiver may not be able to enumerate all books that are neither novels nor poetry collections, but they are presumably not able to enumerate all novels either. A second option is that the agent be able to recognise that the proposition is true when presented with an instance of it. But this is too weak: our agent would recognise whether a book with which they were presented was a novel, but they would also presumably recognise whether a book with which they were presented was neither a novel nor a poetry collection. A third option is that the agent be able to provide an example.¹⁹ But this is unsatisfactory too: the agent can easily give an example of a non-novel non-poetry book, for instance by naming the book that's geographically closest to them and that's neither a novel nor a book of poems. And to say why *Mrs Dalloway* counts as an example of a novel, but the geographically closest non-poetry non-novel book does not count as an example of *other* is precisely to give an account of the intelligibility/unintelligibility distinction. So, if regular and catch-all propositions differ in how intelligible they are to the agent who entertains them, it remains to be outlined precisely what this means.

And even if the intelligibility/unintelligibility distinction could be outlined precisely, it's not clear that it marks the distinction between what we intuitively take to be catch-all and other propositions. For suppose that a theoretical physicist tells me that there are two positive candidates for a theory of fundamental physics: super-symmetry theory and quantum gravity theory. In this situation and knowing that these two theories do not exhaust the space of possibilities, I would entertain *super-symmetry*, *quantum gravity*, and *other*. But, having no training in theoretical physics, I would not have any understanding of what *super-symmetry* or *quantum gravity* say; no more than I would have an understanding of what *other* says—all three propositions would be unintelligible to me. And I take it that, in this example, *super-symmetry* and *quantum gravity* are supposed to be actively entertained, and *other* is supposed to be a catch-all proposition. It follows that unintelligibility, even if it could be precisely theorised, can't be the characteristic feature of catch-all propositions. Many non-catch-all propositions are unintelligible.

In fact, many non-catch-all propositions are also intractable. Our book-receiver might know what novels and poetry collections are but might have no idea how likely they are to receive the one relative to the other, for instance if they have never spoken with their book-gifter about either. Our scientist may know what the two formulated theories say, but

¹⁹ Thanks to an anonymous referee for this suggestion.

may find it difficult to settle on one over the other if it's unclear (as it often is) which the evidence supports. Non-catch-all propositions may even be more intractable than their catch-all counterparts. The book-receiver whose friend mostly talks about experimental literature might find it much easier to determine the probability of receiving a non-poetry non-novel book, than that of receiving for instance a novel. In this kind of case, the ratio of the catch-all proposition to its negation is far more tractable than the ratio of an actively entertained proposition to another.

These considerations shed new light on premise **iii**, the claim that we ought to refrain from modelling intractable propositions. Firstly, our analysis of tractability highlights that it is not a binary affair. Since one can have more or less evidence for a proposition, and a more or less clear sense of a proposition's relation to the evidence, it may be more or less easy to ascertain the probability of a proposition.²⁰ This presents a difficulty for the premise: what, if tractability is gradable, marks the limit between those propositions that may felicitously be employed in epistemic modelling and those that may not? Secondly, the premise no longer establishes what it is intended to establish. Indeed, the defectiveness argument, as part of the argument from awareness growth, is designed to target orthodoxy specifically on the basis of its employment

²⁰These considerations are respectively termed the *weight* ("the gross amount of relevant data available") and *specificity* ("the degree to which the data discriminates the truth of the proposition from that of alternatives") of the evidence in Joyce's classic paper (2005, p. 154). They are sometimes subsumed under the general label of *severe uncertainty*, which Bradley defines as situations in which "the agent is not merely in a state of uncertainty in which, though she doesn't know for sure whether something is true or not, she can assign a probability to it being so. Rather, she is in a state of severe uncertainty, in the sense that, such are the limits on what she knows and can learn, she has no non-arbitrary basis for assigning such a probability" (Bradley, 2017, p. 225). He traces this idea back to Keynes who discusses propositions for which "there is no scientific basis on which to form any calculable probability whatever. We simply do not know" (1937, p. 214). Severe uncertainty has widely been taken to motivate a particular departure from orthodoxy: *imprecise Bayesianism*, on which an agent's epistemic state is represented not by a single probability function but by a set thereof (Bradley, 2019). This is suggestive. If one is inclined to adopt imprecise Bayesianism in the face of severe uncertainty, it seems one might also be inclined to adopt imprecise Bayesianism in the face of what now appears like another instance of severe uncertainty: the sometimes-low tractability of some catch-all propositions. It is curious then that many of the proponents of the defectiveness argument have argued elsewhere for imprecise Bayesianism, for precisely the reason that it handles severe uncertainty (Bradley, 2017, ch. 11; Bradley and Steele, 2014; Romeijn and Roy, 2014). These considerations further fuel my suggestion in fn. 14 that imprecise approaches to catch-all propositions are worthy of interest.

of catch-all propositions. But if there are perfectly tractable catch-all propositions, and quite intractable non-catch-all propositions, the invocation of premise **iii** threatens the specificity of the argument from awareness growth. This brings us to the third and most significant difficulty. Since it is often far from easy to ascertain the relation between a proposition (even if not a catch-all, even if perfectly intelligible) and the evidence, many (perhaps even most) propositions of interest to the Bayesian are to some not-insignificant extent intractable. So, invoking premise **iii** not only fails to specifically prevent the employment of catch-all propositions, it also precludes much of the theorising that Bayesianism is intended, and used, to model. Unless the opponent of orthodoxy is willing to discontinue such theorising, they will have to renounce premise **iii**.

The defectiveness argument thus fails at every step. It follows that, contrary to what seems to be universally believed, the employment of catch-all propositions does not count against endeavours that require it.²¹

3.

Let us zoom back out to the argument from awareness growth:

1. Orthodoxy must accommodate lateral awareness growth.
2. This requires the employment of catch-all propositions.
3. Endeavours that require catch-all propositions must be rejected.
4. So, orthodoxy must be rejected.

This argument is fallacious: as we just saw, premise **3** is false, and an endeavour's requiring catch-all propositions need not count against it. Of course, having established this suffices to defend orthodoxy against the argument from awareness growth. But I want, in this section, to examine the general structure of the argument from awareness growth, quite apart from the truth of its third premise. The argument, I will show, implicitly encodes a view about the relationship between orthodoxy, catch-all propositions, and awareness growth which I shall call the *received view*. This view is obscured by the manifest structure of the argument, but once excavated, can be seen to be false. Establishing this will reinforce my opposition to the

²¹ I have not established that no argument could be offered that would entail that catch-all propositions are problematic, but merely that the argument that is in fact offered for this conclusion does not succeed. If there exists a better argument for the conclusion, the onus is on the critic of catch-all propositions to present it, though the considerations I have surveyed in this section make me doubtful that it can be done.

argument from awareness growth. But, more importantly, it will make visible the actual relationship between orthodoxy, awareness growth, and catch-all propositions, expressed by what I shall call the *revisionary view*. The articulation and adoption of the revisionary view will have important upshots for the issue of how to accommodate awareness growth on orthodoxy—a task I take on in §4.

Let us begin excavating the received view. The argument from awareness growth explicitly encodes in premise 2 the fact that, as we saw in §1, on orthodoxy, the accommodation of lateral awareness growth requires the employment of catch-all propositions. But, implicitly, it also encodes its converse: that it is *in order to* accommodate lateral awareness growth, and not for another reason, that catch-all propositions must be employed. Indeed, the proponents of this argument locate orthodoxy’s vulnerability to the defectiveness argument in its accommodation of awareness growth: they advance an argument *from awareness growth*. It follows that the argument encodes a bidirectional link between catch-all propositions and awareness growth on orthodoxy: not only are catch-all propositions required for the representation of awareness growth, but the representation of awareness growth constitutes *the* impetus for the employment of catch-all propositions.

This is rather striking. Awareness growth (a diachronic phenomenon) takes place against a prior epistemic state (a synchronic phenomenon). So, on the received view, the prior synchronic phenomenon must be modelled with reference to a diachronic phenomenon that could potentially occur thereafter: it is because the agent’s awareness might grow that their initial epistemic state must be modelled with a catch-all proposition; so that there would be a proposition to refine, were a growth of awareness to occur. So we can call the received reason for employing catch-all propositions *proleptic*.²²

We can now articulate the received view. Take an agent whose awareness is laterally limited: there exist possibilities besides the ones actively entertained by the agent. This agent

²²If the received view were right, there would exist another argument from awareness growth against orthodoxy, besides the one that relies on the *defectiveness* of catch-all propositions, that would rely on what we might call their would-be *artificiality*. Indeed, suppose that the reason to employ catch-all propositions in the representation of agents with limited awareness was merely to allow for the potential representation of a potential event of awareness growth. Then, the employment of a catch-all proposition in the representation of the agent with limited awareness would not correspond to a feature of the agent at that time; but instead would be a mere modelling *artefact*—a piece of extra structure included in the model for the purposes of avoiding its possible prospective breakdown. It would follow that the representation of agents with limited awareness would be unsatisfactory in that it would not properly and precisely capture the contours of the phenomenon at hand.

is susceptible to awareness growth; they could become aware of one of these aforementioned possibilities. Were this to happen, its modelling on orthodoxy would require a catch-all proposition; a proposition therefore of a different type than the actively entertained ones. And awareness growth, when it does happen, must be accommodated. It is for that (proleptic) reason that orthodoxy requires modelling the epistemic states of agents with laterally limited awareness using catch-all propositions.

This received view is false. Indeed, as I remarked in §1, orthodoxy entails the Boolean norm. This norm entails that, for any proposition A_i that the agent entertains, the agent must also entertain the complement of A_i against the domain of discourse Ω . And orthodoxy interprets Ω modally: Ω represents the (logically, metaphysically, porically) necessary proposition. So, suppose that an agent is in a state of limited awareness—the disjunction of the propositions they actively entertain is logically stronger than the necessary proposition. For example, our book-receiver entertains the propositions *novel* and *poems*, the disjunction of which entails the necessary proposition. It follows that the agent must also entertain the complement of this disjunction. And this is, by definition, a catch-all proposition. So the fact that orthodoxy must model agents whose awareness is limited with a catch-all proposition follows, not from any consideration about how to model a possible future growth in awareness, but directly from two of its constitutive components: the Boolean norm, and the modal interpretation of the formalism.²³

The systematic employment of catch-all propositions is mandated by core, defining features of orthodoxy. Every time an agent's awareness is laterally limited—every time the propositions they entertain do not disjoin up to the necessary proposition—this agent's awareness state must be represented with a catch-all proposition. So by contrast with what is intimated by the argument from awareness growth, orthodoxy's would-be vulnerability to the supposed defectiveness of catch-all propositions comes, not from the necessity to leave the possibility open for an awareness growth event to be represented, but instead from the orthodox demands of rationality themselves. Were the defectiveness argument convincing, it would be threatening to orthodoxy not simply to the extent that modelling awareness growth is desirable, but much more generally and fundamentally.

It follows that much more hangs on the refutation of the defectiveness argument than it seemed heretofore. For until now, it seemed that the defectiveness argument only posed a challenge to orthodoxy provided that premise 1 held—provided that it was important for

²³ Thanks to Teruji Thomas for helping me simplify what was previously a considerably and unnecessarily more complicated argument.

orthodoxy to accommodate awareness growth. And although there is no defence in print of the view that it is not—that awareness growth exists but need not be accommodated²⁴—the comparatively small volume of literature on the topic, together with the fact that catch-all propositions are only ever mentioned in the context of discussions of awareness growth, suggest that many Bayesians might hold this view. But since the need for catch-all propositions stems from characteristic features of orthodoxy, their employment is widespread, and all orthodox Bayesians should be concerned with the potential threat they pose.

It is a relief then, for orthodoxy, that the defectiveness argument fails. But it is also non-accidental. I showed above that catch-all propositions must be employed whenever the agent’s awareness is limited, because of the Boolean norm and the modal interpretation of the formalism. But this line of reasoning applies not only to propositions intuitively considered to be catch-all propositions, such as in our book case or in the cases drawn from the history of science, but also to more banal cases. For instance, if the book-receiver entertains that the book may be experimental, they must also entertain that it may not be. Or, if an agent entertains that it might rain, they must also entertain that it might not. Here, the propositions *non-experimental* and *non-rain* have the same status as the catch-all propositions in the book and physics cases: they fill out the epistemic space against the propositions initially entertained. But they also have much less of a catch-all flavour; so the distinction between catch-all propositions and not begins to fall away. This intuition is strengthened by noticing that these cases could just as easily have been described with the pairs *conventional/non-conventional* and *dry/non-dry* instead.

This raises the general question of how catch-all propositions feature in orthodoxy, if not proleptically to allow for the modelling of possible lateral growths of awareness. Probabilism states that a rational agent’s epistemic state can be *fully* represented by a probability function—a function defined over a Boolean algebra. In other words, on orthodoxy, the awareness component of epistemic states is entirely captured by the Boolean norm. But the only features of propositions that are recorded by Boolean algebras are their propositional-logical relations. And intelligibility, or whatever feature we take to be characteristic of catch-all propositions, can clearly not be expressed purely in these terms. So, there is no way to express the difference between catch-all and non-catch-all propositions within orthodoxy.²⁵ This sheds

²⁴ For arguments that awareness growth does not exist, or that it is much rarer than what is assumed throughout the literature, see Maher (1995) and Mahtani (2021a).

²⁵ The only way to express this difference without renouncing probabilism is to supply an interpretation to the formalism different from the modal one, on which the trivial proposition represents the disjunction of all actively entertained propositions. This amounts to the kind of relativisation of

light on the relation between what was established in §2 and §3: the Boolean norm is both what mandates the inclusion of (what we intuitively think of as) catch-all propositions beyond cases of prospective awareness growth, and what makes the distinction between catch-all and non-catch-all propositions—and thus the defectiveness argument—impossible.

We now have all the cards in hand to conclude on §2–3. In these sections, I have refuted the argument from awareness growth.²⁶ I have also drawn a picture of the relation between orthodoxy, catch-all propositions, and awareness growth. On orthodoxy, the awareness component of the agent is given by a set of propositions of which the only thing that can be said is their logical relations to each other. Moreover, the credal component is entirely governed by evidential relations between these propositions. So, there is no difference to be internally drawn between non-catch-all and catch-all propositions, either in terms of how they feature in the awareness component, or in terms of how they are to be assigned a probability. There is only a set of propositions, situated in a web of evidential relations, some of which clearer and more instructive than others. With this revisionary view in hand, we are now in a position to provide a suitable account of awareness growth on orthodoxy.

4.

On orthodoxy, lateral awareness growth should be accommodated by refinement. Not only is this mandated by the revisionary view, but it also turns out to be elegant, theoretically inexpensive, and philosophically insightful. So if anything, the requirement to accommodate awareness growth counts in favour, and not against, orthodoxy. Or so I shall argue in this section.

At the start of this paper, I drew a distinction between three types of awareness growth: internal, orthogonal, and lateral. Everyone agrees, I remarked, that internal and orthogonal

the domain of discourse discussed in §1, proposed by almost all endorsers of the argument from awareness growth as an alternative to orthodoxy.

²⁶I have shown that the argument fails on its own terms. But on top of this, the argument is difficult for proponents of anti-orthodoxy to make. For there are obviously cases in which agents actively entertain what have the intuitive flavour of catch-all propositions: the book-receiver who thinks their friend would not gift them anything as banal as a novel or poetry collection, the scientist who thinks the correct theory of their domain of study has likely not yet been articulated. How can the proponent of the argument from awareness growth accommodate these cases? If they represent the agents therein with the help of catch-all propositions, they fall prey to their own plea against them. But if they represent them without, they are not grasping the agents' mental states—what are they doing then?

awareness growth should be represented by *refinement* of the algebra. So, as the book-receiver realises that a poetry collection might be a collection of prose poems or of poems in verse, they must be represented thus: where their initial algebra contains the proposition *poems*, their subsequent one must contain the disjunction of *prose* and *verse*. But this does not fully settle how this agent's awareness states should be represented. It leaves open the question of whether these are the only propositions to be included in the agent's algebra, or whether *other* should also be added besides them. The consensus view of course must be that the agent's initial algebra does not include *other*. For remember that lateral awareness growth had been singled out as uniquely problematic for its requirement of a catch-all proposition. Internal and orthogonal awareness growth by contrast were not thought to be problematic, presumably because they do not involve such a requirement. This makes sense on the received view: if the reason for including a catch-all proposition is only ever to allow for a potential growth in awareness, and since modelling internal and orthogonal growths of awareness as such does not require catch-all propositions, *other* need not be included; indeed since it is supposedly defective, it must not.

But this does not hold on the revisionary view. There, the inclusion of catch-all propositions on orthodoxy is required not for the proleptic reason that it allows for the modelling of potential lateral growths of awareness, but because it is demanded by fundamental features of orthodoxy. This applies to the case of internal awareness growth we are considering here. (Indeed, it applies regardless of whether any growth of awareness takes place.) Since *novel* and *poems* do not exhaust the space of possibilities, the agent's initial algebra must also contain the proposition *other*. Therefore, the complete way to represent this case of internal awareness growth on orthodoxy is as follows. Their initial algebra \mathcal{A} must contain the propositions *poems*, *novels*, and *other*. Their subsequent algebra \mathcal{A}' must also have *novel* and *other*, but where it previously simply had *poems*, it must now have *prose* and *verse*. So, the agent continues to entertain the proposition *poems* as the disjunction of *prose* and *verse*.

On orthodoxy, internal and orthogonal awareness growth should be represented as refinements of an algebra potentially containing a catch-all proposition. Let us then consider the possibility that lateral awareness growth should also be represented in this way. This possibility, applied to the agent who comes to entertain *stories*, would look as follows. Their initial algebra would be such, we have seen, that the atomic propositions are *novel*, *poems*, and *other*. As they become aware of *stories*, their algebra would be refined such that they continue to entertain *novel* and *poems*, but instead of *other*, they now entertain *stories* and *other'*, where these two propositions disjoin up to *other*. So, the proposition *other* ("something other than a novel or a poetry collection") would get split into *stories* and *other'* ("something other than a

novel, a poetry collection, or a book of stories”, that is, from the perspective of the agent who has just become aware of *stories*, “something else still”). On this view, *all* cases of awareness growth, including lateral awareness growth, would be represented in the same way: by refining the algebra. Let us call this the *refinement view*.

My refutations of the defectiveness argument and of the received view make, we have just seen, the refinement view a viable option for how to model awareness growth on orthodoxy. But they do more than this: provided that one is committed to the representation of internal and orthogonal awareness growth by refinement, they entail that one should adopt the refinement view in full generality, including for lateral growths of awareness. For I have argued—this is the revisionary view—that although some propositions, in some contexts, certainly have the intuitive texture of catch-all propositions (the unarticulated theories of our physicists, the mysterious genre of our book-receiver), while others have the intuitive texture of actively entertained ones (general relativity theory, poetry), there is nothing on orthodoxy but a structured set of propositions on which evidential relations are drawn, and therefore, that there is no way to distinguish catch-all from non-catch-all propositions. It follows that what appeared, pre-theoretically, to be a normatively relevant distinction—the distinction between internal and lateral forms of awareness growth—cannot in fact be drawn. For the hallmark of lateral awareness growth was that the new proposition is inconsistent with the old “actively entertained” ones. But if there is no principled way to pick out actively entertained propositions as opposed to catch-all ones, it becomes impossible to express this hallmark, and therefore, to maintain this distinction. So, if one is committed to representing internal awareness growth by refinement, one must be committed to doing the same for lateral awareness growth. This is the refinement view.

My analysis of catch-all propositions and awareness growth on orthodoxy thus entails that all cases of awareness growth must be given a unified treatment. This is theoretically appealing, as even those who advocate for a differentiation between types of awareness growth and of their handling recognise. So Steele and Stefánsson, who insist on the distinction between internal and lateral growths of awareness, call the refinement view “far more elegant” (2021a, p. 1225) than the one they propose, partly on the grounds that it has “the kind of generality [they] seek in a model of awareness growth” (p. 1227). This suggests that, were they to be persuaded by my case against the defectiveness argument and by my defence of the revisionary view, they might happily adopt it.²⁷

²⁷ The refinement view also has another theoretical advantage over the kind of view they adopt. They argue, we have seen, that awareness growth should sometimes be represented by *expansion*, such that the entire space of possibilities changes. This raises difficult questions concerning how to make sense

But more can be said for the refinement view than merely remarking on its elegance and parsimony: it is philosophically illuminating in its handling of awareness growth. To see this, consider the rationale for modelling seemingly internal cases of awareness growth by refinement. The agent in our book case becomes aware that poetry collections might be collections of prose poems or collections of poems in verse. One way to describe what happens as their awareness grows is to say that they come to appreciate a distinction they were previously unable to make; their power of discernment increases. More precisely: what the agent previously took to be one possibility (*poems*), they are now able to cognise more finely as two possibilities (*prose* and *verse*). But the rest of their epistemic landscape remains unchanged; the agent’s cognition of the propositions *novel* and *other* is unaffected by this growth in awareness. So, what has happened is a local precisification of their general awareness landscape. Now, notice that this description of the seemingly internal case is also an illuminating description of what happens in a seemingly lateral case. Whereas at the outset, the agent cannot discriminate between literary genres besides the novel and the poem—all non-novel non-poetic genres are lumped together for them—they come to appreciate a distinction among these genres, between short stories on the one hand, and yet-other genres on the other. But the propositions *novel* and *poem* remain unchanged. So in seemingly lateral cases of awareness growth too, the agent’s general awareness landscape is precisified; in this case at the site of the seemingly catch-all proposition. On this description, awareness growth consists in *increasing discernment*.

As those who model seemingly internal cases of awareness growth by refinement know, this description of awareness growth as increasing discernment is well captured by algebra refinement. For if \mathcal{A}' is a refinement of \mathcal{A} , then by definition, all the propositions in \mathcal{A} are retained in \mathcal{A}' , and new propositions appear as splittings of old ones. This allows us to represent the two features of increasing discernment: the unchanging general structure of the awareness landscape is rendered as the endurance of all the old propositions of \mathcal{A} in \mathcal{A}' ; and the precisification of the landscape is given as the division of some old, coarse-grained propositions of \mathcal{A} into more fine-grained, detailed propositions in \mathcal{A}' .²⁸

of an agent’s continuing to entertain a proposition across awareness growth (Steele and Stefánsson 2021b, Roussos msa). On the refinement view, since the general landscape of possibilities remains constant, and every initially entertained proposition persists, this question does not arise. This is a non-negligible step towards a proper formulation of Reverse Bayesianism: see fns. 8 and 9.

²⁸This precisification can occur in differing degrees of scope. At one extreme, the growth in awareness concerns just one old proposition of \mathcal{A} , in which case the precisification of the awareness landscape has a very local scale. This is what happens in the two examples we have been discussing: the seemingly internal example where the agent becomes able to distinguish prose from verse poetry,

The point can be taken further. As I have characterised them, orthodox Bayesians need not take a stance on the interpretation of the individual elements of Ω ; they need only interpret Ω itself in a modal way. But many orthodox Bayesians have imbued the individual elements $\omega_1, \omega_2, \dots \in \Omega$ with philosophical significance. Most famously, Lewis states that credence functions are defined on “the space *whose points are possible worlds* and whose regions (sets of worlds) are propositions” (1980, p. 267, emphasis added). According to him, the elements $\omega_1, \omega_2, \dots \in \Omega$ represent each of the metaphysically possible worlds, and propositions $A_i \in \mathcal{A}$ are sets thereof. The idea is that each proposition is associated with all the metaphysically possible worlds in which it obtains. This allows us to capture the increasing discernment of agents undergoing awareness growth in an even deeper way. An agent who is unable to tell any difference between some possible worlds and therefore groups them together in a single set, becomes able to discern that some of these possible worlds have a characteristic that the others lack. So, whereas all the possible worlds in *other* seemed identical to the agent, those on which the book is a collection of stories and those on which it is not come to be visible as such. Now of course, the interpretation of individual elements of Ω as possible worlds is controversial and many proponents of orthodoxy reject it.²⁹ So the point here is rather modest: that those orthodox Bayesians who do adopt this kind of interpretation are able, with the refinement view, to capture the increase in discernment of agents undergoing awareness growth in particular detail.

The necessity, on orthodoxy, to model lateral awareness growth by refinement is thus not a reason against it. If anything, I have shown in this section, it is a reason in its favour, for it already features all the theoretical resources necessary to make manifest that, as their

and the seemingly lateral example where the agent becomes able to distinguish short stories from other non-novel non-poetry genres. And at the other extreme, the growth in awareness concerns all the old propositions of \mathcal{A} , in which case the precisification has a global scale. This is what happens in the case I called “orthogonal”, on which the agent becomes aware that all books regardless of genre might be printed as hardbacks or as paperbacks.

²⁹ A major reason for this is that agents sometimes mark a difference between two possibilities which actually correspond to the same metaphysically possible world. For instance, an agent might entertain that a book could be by bell hooks or by Gloria Jean Watkins, as if these were different possibilities, although unbeknownst to the agent, the two names refer to the same person. This is Frege’s Puzzle, discussed in the context of interpreting the Bayesian formalism by Chalmers (2011). It is interesting to note however that the elements of Ω could be given an interpretation other than that of metaphysically possible worlds, on which this problem might not arise. See Mahtani (2021b, forthcoming) for a discussion of some options.

awareness grows, agents become increasingly discerning. Thus we can think of the refinement view not only as mandated by core features of orthodoxy, but also as allowed by it.

5.

Let us wrap up. I began this paper by sketching why many have thought awareness growth to be a reason to reject orthodoxy; I called this the *argument from awareness growth* (§1). I argued that this argument does not succeed: it does not establish that awareness growth constitutes a reason to depart from orthodoxy (§2-3). Instead, I showed, there is an inexpensive and insightful way to model awareness growth on orthodoxy, which I called the *refinement view* (§4). The situation then is the exact opposite of what is usually believed: far from constituting a reason to reject orthodoxy, the need to accommodate awareness growth in fact constitutes a reason to embrace it.

But a crucial step in the argument merits further comment. My case for the refinement view relies on catch-all propositions being employed whenever an agent's awareness is limited. If the propositions actively entertained by the agent do not disjoin up to the necessary proposition—for instance if they only entertain *novel* and *poems*—the rest of the epistemic landscape must be filled out with a catch-all proposition. But there is a possible challenge to this assumption. What about the agent who is genuinely convinced that *novel* and *poems* exhaust the space of possibility—that $novel \vee poems$ is a necessary proposition? Ought this agent not be represented in adequacy with their cognitive state, such that *novel* and *poems* form a partition of Ω ? And if so, how could a lateral growth of awareness in this agent be represented by refinement?

The first thing to note about this challenge is that it does not concern the cases usually discussed in the literature on awareness growth. The kinds of cases discussed there are similar to the physicists' case or the book case as we've been considering them in this paper until now; they are not cases in which the agent mistakenly takes *novel* and *poems*, or the two classical theories, to genuinely exhaust the space of possibilities. The agents in these cases do not make a mistake in taking a contingent proposition to be necessary. Instead, they simply fail to entertain or articulate some alternatives to the propositions that are salient to them. So the challenge at hand does not put the refinement view into question for the cases described in the literature.

Nonetheless, it is a challenge to the refinement view in full generality, and as such merits a closer look. So, what happens when an agent mistakenly takes $novel \vee poems$ to be neces-

sary?³⁰ This agent flouts the disjunction of the Boolean norm and the modal interpretation of the formalism: they fail to properly grasp the logical relations between propositions, or to properly grasp a necessary proposition as such. That such a mistake is possible entails two problems with orthodoxy, well-known since Savage (1967). Firstly, the norm we might call *trivial omniscience*, which is entailed by probabilism, states that agents must have extremal credences in trivial propositions (so, credence 1 in Ω and credence 0 in \emptyset). Together with the orthodox interpretation of the formalism, trivial omniscience requires agents to be absolutely certain in all true necessary propositions. But real agents fail to recognise many true necessary propositions as such, and are thereby deemed irrational by orthodoxy. And secondly, the norm known as *additivity*, which is also entailed by probabilism, states that agents' credences in inconsistent propositions must add up to their credence in their disjunction (so, if A_i and A_j are inconsistent, then $p(A_i) + p(A_j) = p(A_i \vee A_j)$). But real agents may be mistaken about whether two propositions are inconsistent, and thus fail to satisfy the norm of additivity, thereby once again being declared irrational.

The conjunction of the Boolean norm and the orthodox interpretation of the formalism states that the awareness component of agents' epistemic state must be such that they always recognise the relations between propositions, and their modal status. But this is a very strong demand, systematically unmet by real agents. This leads to many problems for orthodoxy: when coupled with trivial omniscience, when coupled with additivity, and as we have just seen above, when coupled with awareness growth. I propose to use the term *problem of logical omniscience* to refer to this root problem. The term of course already exists in the literature, though it is not used evenly. Some use it to refer exclusively to the problem that arises when conjoining this root problem with the credal norm of trivial omniscience. Others also use it also as applied to the norm of additivity. But I am proposing to use the term to refer, not to the unattainable demand of absolute certainty in all necessary propositions, nor to the unattainable demand of additivity, but to what I take to be the root cause of both of these demands on the credal component of agent's epistemic state: the demand on the *awareness* component of the agent's epistemic state that it be logically/modally sound.

Understood this way, the problem of logical omniscience is this—the logical/modal status of individual propositions and of relations between propositions is far from easy to ascertain; and yet orthodoxy requires agents to have ascertained them. There has been disagreement

³⁰Such an agent takes the disjunction of *novel* and *poem* to be a necessary proposition, that is, to exhaust the space of what is a possible option. So it is not the case, for instance, that they ascribe credence 0 to the proposition *other*: they are not certain that the proposition is false, as they might if they had conditionalised on its negation; they do not entertain the proposition in the first place.

about whether to abandon orthodoxy as a response to this problem;³¹ and of course I cannot take a stance on this here. But what matters for the purposes of this paper is this. The challenge to the refinement view we are considering here is not a new problem for orthodoxy. It is merely the application of a well-known problem with orthodoxy—the problem of logical omniscience—to the phenomenon of awareness growth. So, the accommodation of awareness growth poses no *additional* challenge to orthodoxy. Those to whom the challenge appears pressing would therefore do well to concentrate their energies, not on awareness growth, but on the problem of logical omniscience itself.

22nd May 2024

Acknowledgements. I am extremely grateful to Richard Bradley, Liam Kofi Bright, Hylke Jellema, Christian List, Anna Mahtani, Nick Makins, Richard Pettigrew, Jan-Willem Romeijn, Joe Roussos, Orri Stefánsson, and three anonymous referees for insightful comments and probing engagement on previous drafts. For discussion of these issues over the years, I am also grateful to Arif Ahmed, Kevin Blackwell, Catrin Campbell-Moore, Franz Dietrich, Edward Elliot, Leah Henderson, Ko-Hung Kuan, Evan Piermont, Marcus Pivato, Katie Steele, Teruji Thomas, and Aron Vallinder; and to audiences in London, Munich, Cambridge, Groningen and Stockholm. Finally, I am deeply grateful to Nicolas Côté for typesetting the final version of this document when I was too sick to do it.

References

- Bradley, R. (2017). *Decision Theory with a Human Face*. Cambridge University Press.
- Bradley, S. (2019). Imprecise Probabilities. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.
- Bradley, S. and K. Steele (2014). Uncertainty, Learning, and the "Problem" of Dilation. *Erkenntnis* 79(6), 1287–1303.
- Briggs, R. A. and R. Pettigrew (2020). An Accuracy-Dominance Argument for Conditionalization. *Noûs* 54(1), 162–181.

³¹ Classical commentators have tended to think that the problem counts against orthodoxy (Hacking, 1967; Garber, 1983), but more recent interventions have disagreed (Smithies, 2015; Carr, 2022; Canson (ms)). See Titelbaum (2022) for a very nice overview.

- Canson, C. (ms). Why Subjectivism?
- Carr, J. R. (2015). Epistemic Expansions. *Res Philosophica* 92(2), 217–236.
- Carr, J. R. (2022). Why Ideal Epistemology? *Mind* 131, 1131–1162.
- Chalmers, D. J. (2011). Frege’s Puzzle and the Objects of Credence. *Mind* 120(479), 587–635.
- de Finetti, B. (1937). La Préviation: Ses Lois Logiques, Ses Sources Subjectives. *Annales de l’Institut Henri Poincaré* 7, 1–68.
- Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. MIT Press.
- Easwaran, K. (2014). Regularity and Hyperreal Credences. *Philosophical Review* 123(1), 1–41.
- Garber, D. (1983). Old Evidence and Logical Omniscience in Bayesian Confirmation Theory. In J. Earman (Ed.), *Testing Scientific Theories*, pp. 99–132. University of Minnesota Press.
- Glymour, C. (1980). *Theory and Evidence*. Princeton University Press.
- Greaves, H. and D. Wallace (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind* 115(459), 607–632.
- Hacking, I. (1967). Slightly More Realistic Personal Probability. *Philosophy of Science* 34(4), 311–325.
- Hájek, A. (2012). Is Strict Coherence Coherent? *Dialectica* 66(3), 411–424.
- Hájek, A. (ms). Staying Regular.
- Henderson, L., N. D. Goodman, J. B. Tenenbaum, and J. F. Woodward (2010). The Structure and Dynamics of Scientific Theories: A Hierarchical Bayesian Perspective. *Philosophy of Science* 77(2), 172–200.
- Hill, B. (2010). Awareness Dynamics. *Journal of Philosophical Logic* 39(2), 113–137.
- Joyce, J. (1998). A Non-Pragmatic Vindication of Probabilism. *Philosophy of Science* 65(4), 575–603.
- Joyce, J. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives* 19(1), 153–178.

- Karni, E. and M.-L. Vierø (2013). Reverse Bayesianism: A Choice-Based Theory of Growing Awareness. *American Economic Review* 103(7), 2790–2810.
- Karni, E. and M.-L. Vierø (2015). Probabilistic Sophistication and Reverse Bayesianism. *Journal of Risk and Uncertainty* 50(3), 189–208.
- Keynes, J. M. (1937). The General Theory of Employment. *Quarterly Journal of Economics* 51(2), 209–223.
- Laudan, L. (1981). A Confutation of Convergent Realism. *Philosophy of Science* 48(1), 19–49.
- Leitgeb, H. and R. Pettigrew (2010a). An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science* 77(2), 201–235.
- Leitgeb, H. and R. Pettigrew (2010b). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science* 77(2), 236–272.
- Lewis, D. K. (1980). A Subjectivist’s Guide to Objective Chance. In R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability*, pp. 83–132. University of California Press.
- Lewis, D. K. (1999). Why Conditionalize? In *Papers in Metaphysics and Epistemology*, pp. 403–407. Cambridge University Press.
- Maher, P. (1995). Probabilities for New Theories. *Philosophical Studies* 77(1), 103–115.
- Mahtani, A. (2021a). Awareness Growth and Dispositional Attitudes. *Synthese* 198, 8981–8997.
- Mahtani, A. (2021b). Dutch Book and Accuracy Theorems. *Proceedings of the Aristotelian Society* 120(3), 309–327.
- Mahtani, A. (forthcoming). *The Objects of Credence*. Oxford University Press.
- Pettigrew, R. (forthcoming). How Should your Beliefs Change When your Awareness Grows? *Episteme*.
- Poincaré, H. (1905/2018). *Science and Hypothesis*. Bloomsbury. Frappier, M., Smith, A., and Stump, D.J. (Trans.).
- Ramsey, F. P. (1926/1931). Truth and Probability. In *The Foundations of Mathematics*. Kegan Paul.

- Romeijn, J.-W. and O. Roy (2014). Radical Uncertainty: Beyond Probabilistic Models of Belief . *Erkenntnis* 79(6), 1221–1223.
- Roussos, J. (ms). Awareness Growth and Belief Revision.
- Salmon, W. (1990). Rationality and Objectivity in Science or Tom Kuhn Meets Tom Bayes. In C. W. Savage (Ed.), *Scientific Theories*, pp. 14–175. University of Minnesota Press.
- Savage, L. J. (1967). Difficulties in the Theory of Personal Probability. *Philosophy of Science* 34(4), 305–310.
- Shimony, A. (1970). Scientific Inference. In R. Colodny (Ed.), *The Nature and Function of Scientific Theories*, pp. 79–172. University of Pittsburgh Press.
- Smithies, D. (2015). Ideal Rationality and Logical Omniscience. *Synthese* 192(9), 2769–2793.
- Steele, K. and H. O. Stefánsson (2021a). Belief Revision for Growing Awareness. *Mind* 130(520), 1207–1232.
- Steele, K. and H. O. Stefánsson (2021b). *Beyond Uncertainty*. Cambridge University Press.
- Tang, W. H. (2012). Regularity Reformulated. *Episteme* 9(4), 329–343.
- Teller, P. (1973). Conditionalization and Observation. *Synthese* 26(2), 218–258.
- Titelbaum, M. G. (2022). *Fundamentals of Bayesian Epistemology, Vols 1 and 2*. Oxford University Press.
- Wenmackers, S. and J.-W. Romeijn (2016). New Theory About Old Evidence. *Synthese* 193(4), 1225–1250.
- Zynda, L. (1995). Old Evidence and New Theories. *Philosophical Studies* 77(1), 67–95.