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, orthodox Bayesianism cannot account for this phenomenon. As a remedy, two proposals have been put forward: I call them the *expansion* view and the *refinement* view. In the developing literature on this topic, there are no outright defendants of the refinement view: not only are there critics, but even those who endorse it do so half-heartedly. In this paper, I argue resolutely in favour of the refinement view.

0.

Real agents often undergo what is known as *awareness growth*: 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". Importantly, such an epistemic shift does not seem irrational; instead it is clearly is 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.[1] But orthodox Bayesian epistemology does not: this is the *problem of awareness growth*.[2] My aim in this paper is to propose a solution to this problem: I show that there is an attractive and theoretically inexpensive way of integrating awareness growth within the Bayesian framework.

There are two possible ways of modelling awareness growth: by *refinement*, or by *expansion* of the algebra. (I will give definitions of these options later on.) In the expanding literature on the problem of awareness growth, there are no outright proponents of the refinement view: not only are

This is a draft. Comments welcome!

[1] What precisely is meant by "feature" will be discussed in §1.

[2] In the Bayesian philosophy of science literature, this problem is sometimes called the *problem of new theories*. Here, I use the term from formal epistemology because it is more general: it encompasses not only the formulation of new scientific hypotheses, but also more mundane phenomena which are identical in structure from the Bayesian perspective.

there critics, but even those who uphold it do so half-heartedly. In this paper, I argue resolutely in favour of the refinement view. I begin by showing that the problem of awareness growth is indeed an important problem for Bayesians: it arises from several plausible and widely accepted views about the foundations of Bayesian epistemology (§1). I trace out the landscape of possible views on awareness growth—the refinement and the expansion views—and I show that the case for the expansion view and against the refinement view relies on two claims, which I call the *contrivance* claim and the *defectiveness* claim (§2). I examine and refute these two claims in turn (§3–4). Then, I make a positive case for the refinement view (§5), in favour of which I conclude (§6).

1.

Bayesians are characterised by a commitment to the claim that an agent's epistemic state may be represented by a function $p : \mathcal{A} \to [0, 1]$, which assigns a degree of belief, or *credence*, to each proposition that the agent entertains.[3] Thus an agent's epistemic state can be thought of as comprising two elements: what we might call the *awareness* element—the set of propositions of which the agent is aware, represented by the set $\mathcal{A}$; and what we might call the *credal* element—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. This latter, credal element has 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. The orthodox view is 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 $p'(\cdot) = p(\cdot|E)$. In what follows, I will assume these two norms: given the wealth of arguments given in their favour and their correspondingly high degree of popularity, the case for the accommodation of awareness growth within the Bayesian framework will be stronger to the extent that it is consistent with them.[4]

By comparison with the credal element of epistemic states, there has been almost no discussion of the awareness element in philosophy. A few philosophers of science briefly wrote about the issue in the latter half of the twentieth century (Shimony, 1970; Glymour, 1980; Salmon, 1990; Earman, 1992; Maher, 1995), and it was more recently taken up by formal epistemologists and decision theorists (Carr, 2015; Wenmackers and Romeijn, 2016; Bradley, 2017; Mahtani, forthcoming; Steele

[3] Defined as such, Bayesianism is a broad church. As I will show, the problem of awareness growth arises for many but not all Bayesians so defined.

[4] I should remark however that my arguments do not depend sensitively on all features of these norms. It follows that my claims would be consistent with many non-orthodox positions.

and Stefánsson, forthcominga, forthcomingb). But this does not come close to the concern for the credal element of epistemic states. This raises a question: is the relative lack of interest in the awareness element of epistemic states not a sign that there is no significant issue there?

In this section, I show that whether the problem of awareness growth arises depends on the views one adopts about **(1a)** the scope of Bayesian epistemology, **(1b)** the standing of various kinds of idealisation, and **(2)** the nature of the relation between the epistemic and the practical. It turns out that the views which give rise to the problem of awareness growth are increasingly popular in the Bayesian literature, thereby both explaining why awareness has historically received little attention, and establishing the importance of attending to this issue at the present moment. Moreover, although a full defence of these views goes beyond the scope of this paper, I rehearse them in such a way as to highlight to the reader what I think are compelling reasons to hold each of them. It follows that the reader who remains unconvinced that there is a problem with awareness growth in the Bayesian framework will have some work to do to defend his position.

Let me begin with the first relevant issue, **(1a)** the scope of Bayesian epistemology. There are various positions one might adopt on this matter: one might claim that Bayesian epistemology is a theory of how to rationally respond to evidence; or that it is a complete theory of rationality; or, in principle, something else still.[5] Whether the importance of awareness growth for an agent's epistemic life constitutes a challenge to orthodox Bayesianism depends on the position that one adopts about its scope. For if one adopts the position on which Bayesianism is a theory merely of how to respond to evidence, then the importance of awareness growth poses no threat: awareness and its dynamics fall outside of Bayesianism's purview. If however, one adopts a view on which Bayesianism provides a complete account of rationality, at least about matters empirical, then the epistemic importance of awareness growth must be addressed. And whereas older treatments of Bayesianism, mainly in philosophy of science, viewed it as a "comprehensive and unified treatment of inductive reasoning" (Earman 1992, p. 2)[6] or as "a logic of inductive inference" (Howson and Urbach 1989, p. 265)—that is, as a theory of rational response to evidence—newer work, mainly in epistemology, takes it to be a general theory of rationality: thus Titelbaum (2012) views Bayesianism as a theory of "what ideal rationality requires" (p. 24), and Pettigrew (2016a) views it as seeking to provide "the general principles of rational credence" (p. 7)—further examples of Bayesians holding

[5] I am not claiming that any particular set of Bayesian credal norms constitutes a full account, either of evidence acquisition or of rationality in general. Rather, I am claiming that, to the extent that it does not, the goal of Bayesian epistemology as a research programme has not been achieved. Thus when I speak of "scope", I speak of *intended* scope, or target domain of inquiry.

[6] Interestingly, Earman remarks on the "silence" of his contemporaries when it comes to awareness growth, and asserts that the corresponding problem "presents both a[n] interesting challenge and an [...] opportunity for Bayesians" (pp. 112-113). Ultimately, he argues, the accommodation of awareness growth requires a "*non-Bayesian* shift in belief functions" (emphasis mine), which is to say "that no form of conditionalisation [...] will suffice to explain the change" (p. 196).

this view abound.

The epistemic importance of awareness growth has three consequences for Bayesianism intended as a complete theory of rationality. Most basically, Bayesianism must not yield the false conclusion that awareness growth is irrational; in other words, it must *allow* for awareness growth. This can be done easily, by restricting the scope of conditionalisation from the claim that all epistemic changes must take place by conditionalising on new evidence, to the claim that, in the event of evidence acquisition, the agent's epistemic state must change in accordance with conditionalisation. But this is not enough. If Bayesianism is to be a complete theory of rationality, it mustn't simply not rule out awareness growth, it must also *accommodate* it, which is to say, it must also contain descriptions of the phenomenon of awareness growth. And finally, it must also *regulate* awareness growth, that is, it must provide norms of rationality for epistemic changes involving awareness growth. Thus we see that what I had referred to as *the* problem of awareness growth is in fact a series of three imbricated problems: the allowance problem, the accommodation problem, and the regulation problem. In this paper, my primary aim is to show how Bayesianism can and should accommodate awareness growth. I focus on this aim because it is only once awareness growth has been accommodated that the regulation problem can be clearly formulated, and then addressed.[7]

I now turn to the issue **(1b)** of the justification of idealisations, which as we shall see is closely related to the issue **(1a)** of scope. Philosophers who study idealisations often mark a distinction between Aristotelian and Galilean idealisations (Frigg and Hartmann, 2020). Aristotelian idealisations consist in the abstracting away from irrelevant phenomena. Frigg and Hartmann illustrate this with the classical-mechanical model of the planetary system: the aim of this model is to capture the movements of planets, so any property of the planets besides their location at every point in time is idealised away, including for instance, their colour, chemical composition, age, etc. By contrast with Aristotelian idealisations, Galilean idealisations consist in making false assumptions about phenomena that are within the scope of the model or theory, for reasons of tractability. Frigg and Hartmann illustrate such idealisations with the example of a model of motion on an ice rink which assumes the ice to be frictionless, when in fact, there is low but non-zero friction. One important distinction between these two types of idealisation concerns their standing to the theorist. Given a particular domain of inquiry, the theorist is unbothered by Aristotelian idealisations: by definition these idealisations concern what is outside their domain of inquiry. Another way to put this might be that the Arisotelian idealisations delineate the domain of inquiry. But the theorist must consider

[7] The discussion of the regulation problem takes place around the norm labelled *Reverse Bayesianism* by Karni and Vierø (2013, 2015), according to which agents ought not modify their credence in propositions unaffected by awareness growth. The norm is defended by e.g. Wenmackers and Romeijn (2016) and Bradley (2017), and rejected by Carr (2015) and Steele and Stefánsson (forthcominga,b).

Galilean idealisations to be not quite satisfactory, in the sense that a full account of the target domain would require the removal of these idealisations. This is consistent of course with finding them pragmatically satisfactory, that is, perfectly good enough for the purpose for which the theory is deployed.

With this brief primer on idealisations, we are in a position to come to consider idealisations in Bayesianism. There are many things that usual Bayesian models do not capture. For instance, the possibility that I could die in a plane crash is much more vivid to me when I am on a plane than when I am at home with no plane journeys on the horizon. But, assuming that how likely I take a plane crash to be remains constant, these differences in vividness do not feature in a Bayesian model of my epistemic state, for they are usually taken to be rationally irrelevant. This is an example of Aristotelian idealisation: since vividness does not pertain to rationality, and given that we are assuming Bayesianism to aim at giving us a theory of rationality, it simply falls outside of the purview of Bayesianism. But some idealisations, rife within Bayesianism, have a more controversial status. Prominent examples, roughly defined, include:

> *Logical omniscience.* Agents are always certain that all logical truths are true and all logical falsehoods are false.
>
> *Computational omniscience.* Agents always recognise the deductive-logical relations between propositions they consider.
>
> *Evidential omniscience.* Agents are always certain of their evidence.

The status of these idealisations has changed as the scope of Bayesianism has widened: while at least some of them were arguably Aristotelian idealisations when Bayesians were concerned with providing a theory of belief change in response to evidence, they can no longer be so—if Bayesianism is to supply a complete theory of rationality, these decidedly fall within its scope. This has spurred a recent trend towards de-idealisation, under the banner of "bounded rationality". Prominent recent examples of de-idealisation relating to logical and computational omniscience include Pettigrew (forthcoming), building on Hacking (1967); and of de-idealisation relating to evidential omniscience include Bronfman (2014), Schoenfield (2017), and Gallow (forthcoming).

How does this relate to awareness? Many Bayesians with the aim of giving a general theory of rationality have have deployed their classic move to handle questions of awareness: idealisation. This is explicit for instance in Pettigrew (2016b) and Hájek and Lin (2017). Specifically, these Bayesians have assumed:

> *Propositional omniscience.* Agents are aware of all possible propositions throughout their

epistemic life.[8]

Bayesianism looks as follows with such an idealisation. Agents are assumed to start out their epistemic life with a credence function defined over a massive algebra of propositions, containing all possible propositions, and possessing no evidence. At this stage, agents are what David Lewis calls "superbabies".[9] Then, Bayesians describe the way in which agents' confidence in these propositions changes as they acquire new evidence.[10] So, if propositional omniscience is an idealisation that can be justified, the need to allow, accommodate, and regulate awareness growth disappears. But this move does not look promising. Since awareness growth is clearly an important aspect of epistemic thriving, it must feature in a complete account of rationality. Indeed it would be odd to claim that Einstein was irrational prior to the formulation of his theories of relativity, or for that matter that their formulation was an instance of irrationality. Thus propositional omniscience cannot be an Aristotelian idealisation, in the way that it might have been if the scope of Bayesianism had been restricted to learning from evidence. No, it is at best a Galilean idealisation. This means that, while the idealisation may be pragmatically satisfactory in particular cases, it cannot be fully satisfactory in general. More precisely, it prevents Bayesianism from being a general theory of rationality, and as a result, there are particular cases within it scope that Bayesianism would be impotent to handle—such as for instance, the case of Einstein just described. Propositional omniscience might be harmless (and useful!) for a number of particular cases, but it must be rejected in general. And without this idealisation, the need for accommodating awareness growth returns.

This brings us to the last, more separate issue, **(2)** the nature of the relation between the epistemic and the practical. Perhaps besides allowing us to capture the fact that epistemic attitudes come in degrees, the greatest strength of Bayesian epistemology is that it is the only epistemological framework on which the relation between epistemic states and actions is theorised. This theorising is done under the umbrella of decision theory. There are many candidate decision theories, the most prominent of which is expected utility theory, which I shall use to illustrate my points.[11] Consider the following decision situation, the *umbrella case*. An agent is about to leave the house, and is faced with two practical options: taking an umbrella, or not taking an umbrella. The feature of the world that is relevant to this decision is whether it is raining. Let us assume that the utility

[8] The reader might feel that this is a rather vague definition; indeed it is not obvious what is meant by *all possible* propositions. I am not particularly bothered by this lack of clarity because, as I will argue, I think this idealisation is ultimately unwarranted. But I want to note that a defender of propositional omniscience would be tasked with the precise formulation of its statement.

[9] This term is attributed to Lewis by Hájek (ms).

[10] The assumption of propositional omniscience does not mandate acceptance that the only way to change one's credences is by acquiring new evidence. For instance, it is perfectly consistent with the view that one can rationally forget evidence (Titelbaum, 2012).

[11] Nothing I will say hangs sensitively on this choice. For instance, all my claims are consistent with risk-weighted expected utility theory Buchak (2013) or indeed Bradley's own view (2017).

the agent derives from each possible outcome is given by the table below:

	Raining	Not Raining
Umbrella	2	1
No Umbrella	-2	5

Expected utility theory tells us that $U \succ \neg U$ if and only if $p(\textit{rain}) < p(\neg \textit{rain})$. The phrase on the left-hand-side is typically interpreted practically, as a claim about taking the umbrella versus not; and the phrase on the right-hand-side is typically interpreted epistemically, as a claim about the agent's credences in whether it will rain. But how is each side, and the relation between them, interpreted exactly?

Possible views on this matter may be divided into two types, which I shall call *constitutivism* and *non-constitutivism*. On the first type of view, the relation between the epistemic and the practical is one of constitution, such that the epistemic is practically constituted. In the Bayesian context, the most prominent constitutivist view of the nature of credences is *dispositionalism*, the view according to which having particular credences just is having particular dispositions to act, given fixed desires.[12] On this view, to say that the agent has a higher credence in $\neg$*rain* than in *rain* is simply to say that the agent is disposed to leave the house with an umbrella, given the utilities stipulated in the table. By contrast, the most prominent non-constitutivist views of credences are versions of *mentalism*, on which the relation between the epistemic and the practical is not one of constitution, but of rationalisation.[13] (I talk of versions of mentalism in the plural because, as it shall become relevant in §3, there are several of them.) On this kind of view, to say that the agent has a higher credence in $\neg$*rain* than in *rain* is to point to a feature of their psyche—to say maybe that the agent is confident that it won't rain—and the claim of expected utility theory is interpreted as saying that, given the utilities provided in the table, this mental state makes it rational to leave the house without an umbrella.[14]

To see how the constitutivism/non-constitutivism debate is relevant to the question of awareness growth, consider the following *book case*, which I will use as a running illustration throughout the rest of this paper. An agent has received a book as a present, of which they considers the genre: they are actively considering that the unwrapped book might be a novel or a collection of poems. Furthermore, they have, let us assume, never heard of short stories. So, our agent is aware of the

[12] This view is sometimes put in terms of dispositions to bet, but as Ramsey puts it, "all life is in a sense betting" (1926).

[13] Some think that this relation is not merely one of rationalisation, but that it is also one of causation and/or explanation: the agent's mental states causes them to perform particular actions in particular contexts; it explains why the agent acts as they do in particular contexts, etc. For such views, see Dietrich and List (2016) and Okasha (2016).

[14] Note that the combination of $U \prec \neg U$ and $p(\textit{rain}) < p(\neg \textit{rain})$ is irrational on a non-constitutivist view, and metaphysically impossible on a constitutivist view.

propositions we might denote as *novel* and *poems*, but unaware of *stories*. Now, whatever the precise nature of what distinguishes *stories* from *novel* and *poems* is in this case, it is a mental feature of the agent. So since an agent's epistemic state is partly constituted by their awareness state (together with their credal state), this entails that the agent's epistemic state is partly constituted by a mental phenomenon, and thus, that constitutivism is false. In other words, the pressure to allow, accommodate, and regulate unawareness and awareness growth can only arise against a non-constitutivist background.[15] And it is well-known that early Bayesians were devoted constitutivists, in particular Ramsey (1926/1931) and de Finetti (1974/1991). However, the numerous problems with constitutivism, documented by Eriksson and Hájek (2007), have led to a mentalist turn, not just in philosophy but also in economics (Dietrich and List, 2016; Okasha, 2016). And given that, on a mentalist account of epistemic states, a distinction exists between propositions of which the agent is aware and propositions of which they are unaware, this mentalist turn is accompanied by the emergence of the problem of awareness growth.

Let me summarise. There has been a shift in the Bayesian literature on the scope of Bayesianism (and thus on the standing of various forms of idealisation) and on the nature of the relation between the epistemic and the practical. As I have shown in this section, the problem of awareness growth does not arise on the older views, on which Bayesianism is a theory of learning from evidence and epistemic states are practically constituted. It does however arise on the newer views, on which Bayesianism is a complete theory of rationality, and epistemic states are mentally (not practically) constituted. As I have shown, these newer views are extremely widespread in contemporary Bayesian epistemology. This does not entail, of course, that they are true, but it does lend some plausibility to them; or put another way, it shows that the awareness sceptic has non-negligible work to do. In any case, I shall assume them here.

2.

My aim in this section is to bring out the salient features of the debate surrounding the accommodation of awareness growth, and to formulate my thesis within the landscape of this debate. To

[15] This kind of consideration is closely related to the argument that Mahtani (forthcoming) provides for the impossibility of unawareness on the dispositional account of credences. Her argument is roughly this. The agent of the book case has a number of relevant dispositions. For instance, they would bet on the proposition that the book is a book of stories, were they presented with that bet. In exactly the same way, they would bet on all (mundane) propositions, were they presented with a bet on them. (Note that this does not assume that for each proposition, there is a unique amount of money that they would bet. More generally, it does not assume anything about what kind of dispositions the agent has. It merely assumes that they have dispositions.) It follows that they have epistemic attitudes towards the proposition that the book is a book of short stories, and in fact towards all propositions: there is no proposition of which they are unaware.

begin, it will be useful to distinguish three types of awareness growth as thought of pre-theoretically. Consider again the agent who knows that they will receive a book of fiction as a present, and who considers that it might be a novel, or a collection of poems. There are three types of awareness growth that they might undergo. Firstly, they might come to consider that the book might be a paperback or a hardback. They would then consider four possibilities: paperback novel, hardback novel, paperback poetry, and hardback poetry. I call this *orthogonal* awareness growth: the newly considered 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 consider three possibilities: novel, prose poems, and verse poems. I call this *internal* awareness growth: the newly considered 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 consider three possibilities: novel, poems, and stories. I call this *lateral* awareness growth: the newly considered proposition is inconsistent with the old ones.

How can the Bayesian model these epistemic events? As mentioned above, in Bayesianism, the propositions that an agent has epistemic attitudes towards are represented by a Boolean algebra $\mathcal{A}$. There are two ways, in principle, to enlarge this algebra to another $\mathcal{A}'$. The first is by *refinement*: $\mathcal{A}'$ is a fine-graining of $\mathcal{A}$. The second is by *expansion*: $\mathcal{A}$ is replaced by a new algebra $\mathcal{A}'$ on a larger underlying space. These two operations are illustrated in the figure below.

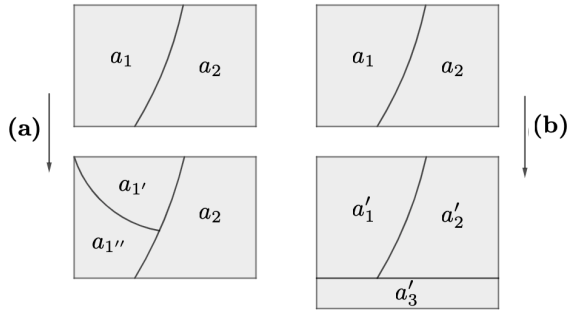

Refinement (a) and expansion (b) of an algebra of propositions

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 consider 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 $\mathcal{A}'$. However, there is a dispute about how to model lateral growths of awareness. Should it be modelled by refinement (the *refinement view*) or by expansion (the *expansion view*)?

An early and influential discussion of awareness growth can be found in the work of Earman

(1992).[16] Earman assumes that awareness growth would have to be modelled by refinement; that is, he does not even entertain the expansion view. He then recognises that the refinement view mandates the inclusion of what Shimony (1970) had called a *catch-all* proposition in the agent's algebra. Roughly, a catch-all proposition is a proposition which expresses something like "a possibility beyond those I have actively considered"; more precisely, it is the complement of the union of all "actively considered" propositions.[17] Why does the refinement view mandate the inclusion of such a proposition? Well, on the refinement view, awareness growth consists in the fine-graining of a previously entertained proposition. But the kind of awareness growth we are considering, lateral awareness growth, is such that the newly considered proposition is inconsistent with the old actively considered ones. So, there must have been an old proposition besides the actively considered ones to be fine-grained: the catch-all proposition. Having recognised that the refinement view requires the inclusion of a catch-all proposition, Earman raises a concern about such an inclusion, which he takes to be fatal, and he concludes that Bayesians cannot model awareness growth. It is only two decades later, in the formal epistemology literature, that we find another take on the problem of awareness growth. Wenmackers and Romeijn (2016) continue to assume that awareness growth must happen by refinement, but do not take the aforementioned concern to be fatal to the view. Instead, they put forward a positive proposal to address this concern. By contrast, Bradley (2017) and Steele and Stefánsson (forthcominga,b) do take concerns about the refinement view to be fatal, but instead propose and argue for the expansion view, and as such hold that the problem of awareness growth can be solved.

We see that there is significant disagreement between authors in the literature. However, everyone agrees on one point: there is an important concern with the refinement view. What is it? The problem is that, unlike other propositions, it is uniquely difficult to assign a rational prior to the catch-all proposition.[18] Earman (1992) claims that assigning such a prior would have to be done in a way that is "arational" (p. 197). In a similar vein, 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?" (p. 255). Wenmackers and Romeijn (2016) write that "unlike the other hypotheses, it is not produced by a scientific theory, but rather it results from a meta-theory" (p. 1333), and they conclude that "since the catch-all is not

[16] Relevant earlier work includes Glymour (1980) and Salmon (1990).

[17] I have put "actively considered" in scare quotes because, although the distinction between actively considered propositions and catch-all propositions might be intuitively appealing, I will in fact argue in §4 that there is no normatively relevant such distinction. Nonetheless it is helpful to speak in this way for ease of exposition.

[18] This is not quite right: Steele and Stefánsson's dissatisfaction with catch-all propositions is slightly different to the one that has been rehearsed in the literature until their intervention. I will present and discuss their dissatisfaction in §4.

based on a scientific theory, the usual 'arational' considerations [...] for assigning it a prior, namely by comparing it to hypotheses produced by other theories, do not come into play here" (p. 1234). According to Henderson et al. (2010), the refinement view "is an unsatisfactory solution since there is no particularly principled way to decide how much initial probability should be assigned to the catchall" (p. 190). The consensus is then that the catch-all proposition is defective: it is uniquely difficult to know what credence to assign to it. I will call this the *defectiveness claim*.

And besides a belief in the defectiveness claim, we can discern another belief, held in common between all these authors. They hold not only that the inclusion of a catch-all is mandated by the refinement view, but the converse, too: that the refinement view is the only reason to include a catch-all proposition in the agent's algebra. This claim plays a crucial role in the argumentative strategies of Shimony, Bradley, and Steele and Stefánsson: they reject the refinement view because it requires the inclusion of a catch-all proposition, and instead defend the expansion view, which they think does not. This dialectical move would not be available to them if they did not think that the inclusion of catch-all propositions was idiosyncratic of the refinement view; that is, if catch-all propositions were to be included on the expansion view, and more generally for reasons other than to formulate the refinement view. Furthermore, although Wenmackers and Romeijn's arguments do not rely sensitively on this claim, it seems to be operative in the background nonetheless. Indeed, an important part of their defence of the refinement view consists in the defence of the inclusion of catch-all hypotheses. But if they thought that the inclusion of a catch-all hypothesis was mandated for reasons independent of the refinement view, there would be no need for such a defence. Thus it is widely held that, not only are catch-all propositions defective, but their inclusion in an agent's algebra is contrived to uphold the refinement view. I shall call this second claim the *contrivance claim*.

In the next two sections, I seek to refute these two claims. In §3, I argue that the inclusion of a catch-all hypothesis is mandatory for reasons independent of awareness growth. So, the contrivance claim is false. And in §4, I argue that there is no difference between the catch-all proposition and the other propositions that would entail the impossibility of assigning a prior to the former and not to the latter. So, the defectiveness claim is false.

3.

To argue against the contrivance claim, I must begin by presenting the formal framework. According to probabilism (which, as I flagged in §1, I will be assuming throughout this paper),[19] the

[19] I should remark here that this assumption is shared by all writers on awareness growth, with the exception of Roussos (ms). I preserve this assumption because (1) it encodes a commitment to classical logic, which I do not wish to give up, especially given the existence of an appealing alternative; and because (2) as I

set $\mathcal{A}$ which represents the awareness element of the agent's epistemic state must form a Boolean algebra, which is to say that it must contain the propositions we might call ***trivial***, represented as Ω and $\emptyset$; and propositions we might call ***non-trivial***, represented as $A_1, A_2, ...$, which must be closed under union and negation. And, again according to probabilism, the credal element of the agent's epistemic state, which is represented by the shape of a credence function $p : \mathcal{A} \to [0, 1]$, must be such that $p(A_i) + p(A_j) = p(A_i \vee A_j)$ for all mutually inconsistent propositions $A_i, A_j \in \mathcal{A}$; and it must be such that $p(\Omega) = 1$ and $p(\emptyset) = 0$. Let us mark the difference between ***extremal*** credences (credence 1 and credence 0), and ***non-extremal*** credences (other credal values between 0 and 1).

A full interpretation of this formalism is not necessary for our purposes; we do however need an interpretation of the trivial/non-trivial distinction, and of the extremal/non-extremal distinction. The latter distinction is typically interpreted as the distinction between asbolute certainty and less-than-absolute-certainty. By contrast the former distinction has been the matter of (mild) dispute in the Bayesian literature. The two most widely discussed contenders state that the distinction marks (i) the distinction between logically necessary and logically contingent propositions; or (ii) the distinction between the metaphysically necessary and the metaphysically contingent (Hájek, 2012; Easwaran, 2014). I argue elsewhere for a third option, namely that it marks (iii) the distinction between the determinable ***a priori*** and the determinable ***a posteriori*** ([redacted]).[20] While the differences between these interpretations are important for various reasons and purposes, they will not matter for the purpose of arguing against the contrivance claim. So, I shall use the terms ***necessary*** and ***contingent*** as an umbrella term to encapsulate the three. I will call the interpretation of the Bayesian formalism on which the trivial/non-trivial distinction is interpreted as the necessary/contingent distinctions and the extremal/non-extremal distinction is interpreted as the absolute-certainty/less-than-absolute-certainty the ***usual*** interpretation, after its widespread popularity.

Many if not all Bayesians have been convinced that agents ought not be absolutely certain of contingent propositions about the unobserved—propositions such as "the sun will rise tomorrow", or to return to our thread example, "the book is a book of poems". I shall call this the ***plausible thought***. On the usual interpretation of the Bayesian formalism, and in contexts where the algebra is finite,[21] this thought can be captured by the following norm of rationality:

> ***Humility.*** $p(A_i) \neq 1$ and $p(A_i) \neq 0$, for all $A_i \in \mathcal{A}$ such that $A_i \neq \Omega, \emptyset$, and $A_i \notin \mathcal{E}$, where $\mathcal{E}$ is the subset of $\mathcal{A}$ containing the propositions that constitute evidence for the

remarked in §1, it is almost universally assumed by Bayesians, on the basis of a range of strong arguments.

[20] 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 (forthcomingb).

[21] This restriction serves to circumvent what might be called the ***too-many-worlds*** objection, the classic discussion of which can be found in Hájek (2012,ms) and Easwaran (2014).

agent.[22]

Humility entails the negation of the contrivance claim as follows. In our book case, the proposition *novel* $\vee$ *poems* is neither evidential—the agent has not unwrapped the book; nor necessary—it is neither logically nor metaphysically necessary, nor determinable *a priori* that a book be either a novel or a poetry collection. So, humility entails that the agent must have a non-extremal credence x in *poems* $\vee$ *novel*. And if the agent has a non-extremal credence x in *novel* $\vee$ *poems*, probabilism dictates that the agent ought to have a non-extremal credence of $1 - x$ in $\neg$(*novel* $\vee$ *poems*). But $\neg$(*novel* $\vee$ *poems*) is, of course, a catch-all proposition: its content can be articulated as "the book I have received is a book, which is neither a novel nor a poetry collection". This argument generalises to all cases of less-than-full awareness: in such cases, agents must be modelled as having a credence in a catch-all proposition, and the contrivance claim is false.

This argument could be resisted by denying the plausible thought; that is, by allowing agents to be absolutely certain of contingent propositions about the unobserved. But this is a very unattractive strategy, and all the more so for the Bayesian, given that Bayesianism is often harnessed to express a response to a certain type of inductive scepticism: "we cannot be absolutely certain that the sun will rise tomorrow," the Bayesian concedes, "but we can have a justified credence that it will." The first conjunct of such a line would not be available to the opponent of the plausible thought; and this would require a repositioning of Bayesianism in relation to the problem of induction.[23] Furthermore, the plausible thought is accepted by all proponents of the contrivance claim; thus Bradley asks: "what should [the agent] do if they are conscious of the possibility *(as*

[22] This norm is weaker than *regularity*, the norm usually at the centre of discussions around the plausible thought. Indeed regularity insists that *all* non-trivial propositions must be assigned non-extremal credences, whereas humility restricts this to non-evidential non-trivial propositions. I stick with humility because regularity clashes with conditionalisation when unrestricted (conditionalisation requires the assignment of extremal credences to evidential propositions), and I believe restrictions to superbaby credences to be theoretically infelicitous, as per my arguments of §1. I discuss this in greater detail in [redacted].

[23] The only remotely promising way to reject the plausible thought, I think, would be to point to cases in which the agent is mistakenly certain that a proposition which is in fact contingent is necessary. One might imagine, for instance, an agent who is absolutely certain that *novel* $\vee$ *poems* is logically or metaphysically necessary, or that it is determinable *a priori*. It could be argued that an agent thusly mistaken ought to assign credence 1 to the propositions in question (in some sense of "ought"). This kind of case is analogous to cases in which an agent is mistaken about their evidence; they might, for instance, be absolutely certain that their evidence entails that the external world exists, when they are in fact deceived by an evil daemon. Such an agent is, in the words of [redacted], "epistemically alienated": presuming that they ought to (and do) conditionalise on the proposition of which they are absolutely certain, they can never retrieve the truth (private correspondence). I leave this kind of case for future work because it is marginal to the question of awareness growth (paradigmatic cases of awareness growth—the book case, the formulation of new scientific theories, etc.—are not cases in which agents are mistaken in this way), and because it requires a unified treatment with other kinds of mistakes, such as mistakes about one's evidence, which go beyond the scope of this paper.

they should be) that they may not be aware of all relevant prospects?" (p. 254, emphasis mine). So, I will take the plausible thought for granted here; and since it is captured by humility on the usual interpretation of the formalism, I will assume that, on this interpretation, humility holds.[24] The falsity of the contrivance claim follows.

However, the line above suggests another way to reject my argument against the contrivance claim: the rejection of the usual interpretation of the Bayesian formalism. Indeed, this is the strategy that both Bradley (2017), and Steele and Stefánsson (forthcominga) adopt. Their first step is the appeal to what Steele and Stefánsson call an "awareness context", which is characterised by the set of all basic propositions of which the agent is aware. Bradley (2017) characterises propositions of which the agent is aware as those which are "available to the agent's consciousness at the time at which they are deliberating on it" (p. 253). His illustrations of unawareness include cases where one "fails to recall [a particular proposition] at a particular time because it slips one's attention" (p. 253), and where one's "attention shifts to other things" (p. 257). In a similar vein, Steele and Stefánsson's examples feature agents who, when pondering what to do, consider a set of possibilities, and later "realise" that other possibilities are also relevant. For instance, a cinema-goer "realises" when rethinking whether to go to the cinema that the film could be a thriller. It is implicit that the agent is not discovering there and then that films can be thrillers; rather, she had simply failed to pay attention to this option. This suggests a particular account of awareness, which one might call the *conscious awareness* account: an agent is aware of a proposition when she is consciously entertaining it, perhaps in the context of active deliberation.

This appeal to an awareness context allows for a reinterpretation of the Bayesian formalism, which I shall call the *contextualist* interpretation. On this interpretation, the trivial/non-trivial distinction is indexed to the awareness context, such that trivial propositions are true/false at all possibilities of which the agent is aware, whereas non-trivial propositions are true at some but not all such possibilities. It follows that whether a proposition counts as trivial depends on the awareness context. To illustrate: the agent in our book case is initially aware of two propositions, *novel* and *poems*, and the proposition *novel* $\vee$ *poems* is a trivial proposition in this initial awareness context. Then, upon becoming aware that the book could be a book of short stories, the agent becomes aware of an additional basic proposition, *stories*, thereby ending up in a new awareness context, in which *novel* $\vee$ *poems* is not a trivial proposition but *novel* $\vee$ *poems* $\vee$ *stories* is. Accordingly, the extremal/non-extremal distinction is reinterpreted as indexed to an awareness context, such that extremal credences do not represent absolute certainty, but certainty in the relevant awareness context. So, probabilism on this interpretation entails that the agent in the book case must have credence 1 in *novel* $\vee$ *poems* in the initial awareness context, which is to say that they must be

[24] I defend the plausible thought in the form of humility in [redacted].

context-certain of this proposition, and upon awareness growth, must cease to be context-certain of *novel* $\vee$ *poems* and become context-certain of *novel* $\vee$ *poems* $\vee$ *stories*.

This reinterpretive strategy can be used to block my argument against the contrivance claim whilst preserving the plausible thought, in the following way. My argument relied on the impermissibility of absolute certainty in propositions like *novel* $\vee$ *poems*, together with the identification of credence 1 with absolute certainty, to argue that a rational agent must have a non-trivial credence in a catch-all proposition, in this case, $\neg$(*novel* $\vee$ *poems*). But because, on the contextualist interpretation of the Bayesian formalism, absolute certainty is divorced from credence 1, the plausible thought does not entail that the agent must have a non-trivial credence in *novel* $\vee$ *poems*, and so, does not entail that the agent must have a non-extremal credence in a catch-all proposition. (In fact, on the contextualist interpretation, the agent must have a trivial credence in *novel* $\vee$ *poems*, and thus may not have a non-trivial credence in a catch-all proposition.) What this shows is that my argument depends on the infelicity of the contextual interpretation of the Bayesian formalism, and the corresponding felicity of the usual interpretation. (These are not exactly exhaustive, which I will discuss shortly.)

Thus I turn to the evaluation of interpretations of the Bayesian formalism. As I outlined above, on the contextualist interpretation of Bradley, Steele, and Stefánsson, one is committed to the claim that our book-receiver ought to have credence 1 in *novel* $\vee$ *poems* in the initial awareness context. Assuming expected utility theory, this entails that it would be rational for our book-receiver to bet everything they have—to bet their life!—on this proposition;[25] all the while, to cite Bradley again, "they are conscious of the possibility (as they should be) that they may not be aware of all relevant prospects" (p. 254). This pragmatic consequence of the interpretation is clearly undesirable. Of course, it could be avoided by severing the link between credence and rational action, and renouncing decision theory. But I take it that no Bayesian would be enticed by such a move; indeed as I have already explained, the capacity of Bayesian epistemology to theorise the link between the epistemic and the practical is one of the (if not *the*) most significant considerations in its favour. And, I shall argue, the fact of this undesirable consequence reveals a conceptual problem at the core of the contextualist interpretation; namely, that it is inconsistent with the thought that (mentally interpreted) credences rationalise action.

In §1, I situated the present discussion within the mentalist family of accounts of credences, on which credences are mental states. And as I explained, the formal relationship between the epistemic and the practical (which is given by decision theory) is interpreted on mentalist views as one of rationalisation: given fixed desires, epistemic states rationalise actions. In what follows, I will argue that this interpretation of the relation as one of rationalisation is not available on

[25] As flagged in fn. 11, this entailment holds on all standard decision theories, including Bradley's own (2017).

all mentalist views of credences; more specifically, it is not available on the conscious awareness account of credences that underpins the contextualist interpretation of the formalism, but it is available on another, widespread account of credences, the representational account, which is naturally associated with the usual interpretation of the formalism, and on which my argument against the contrivance claim goes through. It follows that, if we are to preserve the Bayesian insight that, given fixed desires, credences determine the rationality of our actions, we ought to model less-than-fully aware agents as having a non-trivial credence in a catch-all proposition, and the contrivance claim is false.

Let's start by saying more about this rationalisation relation, with the help of the umbrella case of §1. Let us assume that the agent has a high credence that it is not raining, and let us assume also that it is in fact raining. Is it rational for the agent to take an umbrella? Well, there is one sense in which it is: they prefer having an umbrella if it's raining, and it's raining. So, given the way the world is, their desires would be most satisfied if they take an umbrella. However, there is another sense in which it would be irrational for them to take an umbrella: they are very confident, after all, that doing so would result in a state of affairs they wish to avoid. The first of these kinds of ought is epistemic-independent: what the agent ought to do in this sense can be determined without reference to their epistemic state. The second one is epistemic-dependent: in order to determine what they ought to do in this second sense, one must consult their credences. Unsurprisingly, it is the second, epistemic-dependent sense of ought that Bayesians are interested in; and as such, that constitutes our target. What we as Bayesians need is an account of credence that is consistent with the thought that what would make it irrational for them to take an umbrella is that they take the world to be such that they would be unhappy if they did. In other words, we need an account of credences on which credences constitute the agent's epistemic perspective on the world, from which possible actions are to be rationally evaluated.

To evaluate the conscious awareness account in light of this requirement, let us consider two examples. (1) Consider a thirsty agent who picks up a glass of water whilst thinking about something else. By assumption, the agent does not consciously entertain the propositions that there is water in the glass, that water quenches thirst, that the glass must be picked up to be drunk from, etc. But it does not follow that he does not believe these propositions. Indeed, he may have a (stored) representation of the world as described by them, even if his attention is not directed towards it, and this is what explains the rationality of his action. (2) Consider a thermostat, which represents the world as being in a particular way (the room as having a particular temperature), and is disposed to act on this basis (is disposed to, e.g. turn on the heating if the temperature is below 19°C). One way to describe this thermostat, which is widespread among decision theorists and which Bradley himself endorses (p. 63), is that the thermostat has beliefs about the room temperature; beliefs

which rationalise its heater-toggling actions.[26]

Part of the epistemologist's job is to isolate the parts of agents' mental landscapes which are epistemically relevant. And according to the proponents of the conscious awareness account, only the components of this landscape that are held in conscious awareness (or that are attended to, or that are actively entertained) can be epistemically relevant; that is, can count as credences.[27] What the above examples show is that this putative necessary condition on credences is inconsistent with the view that credences rationalise action. The first example shows that the only way to understand the water's drinking as rational is to accept that the agent has a variety of beliefs that are not presently vivid to him. Indeed a great number of beliefs are relevant to the rationality of our actions—many of which often, if not always, fail to be consciously attended to at the time of action. This is illustrated in a more drastic way by the second example. On this example, it makes sense to evaluate the rationality of an agent's actions in light of its beliefs, even if this agent does not merely contingently fail to be conscious of the relevant beliefs, but necessarily so fails, for it (presumably) could not entertain anything consciously. If the vast majority of our actions are to be rationally evaluable, it cannot be the case that epistemic attitudes must be consciously entertained.

Instead, our epistemic attitudes must be something like our (not necessarily conscious) representation of the world—how, in the phrase I have been using, the agent "takes the world to be". To express this, Bayesians sometimes invoke Ramsey's claim that epistemic attitudes are the "map by which we steer":[28] the agent's representation of the world which functions as the measure of the rationality of their actions.[29] The picture is suggestive: just like a driver might not consult their map when travelling through in well-known places but pay close attention to it in foreign lands; so might an agent's epistemic state reside below the surface of her consciousness when she executes tasks such as drinking some water, but emerge when engaged in scientific inquiry. Indeed I suspect that closely related considerations led Bradley, Steele, and Stefánsson to their views. They focus, in their work in awareness growth, on agents in the process of deliberating; for instance, an agent is wondering whether to go to the cinema. Since it makes sense that agents would deliberate only in case of harder decisions, and since it makes sense that theorisers would focus on the harder, more prominent questions, it makes sense that the more mundane cases would be missed. But once we

[26] This example somewhat stretches the usual meanings of the terms "belief", "desire", "mental", "agent", and so on; but despite this it is very popular among decision theorists.

[27] The converse is clearly false: a great proportion of our conscious life is normatively irrelevant—how vivid various possibilities are to an agent, which she hopes for, and, in fact, much of the rich texture of our conscious experience is epistemically moot. Thus conscious awareness of a mental state is not sufficient for that state to count as epistemic.

[28] Although the phrase comes from Ramsey (1926), he was not a representationalist about credence but a dispositionalist—see §1.

[29] It may also play additional roles such as *causing* the agent to act in particular ways, or *explaining* why the agent acts in particular ways. See fn. 11.

turn our attention to them, we see that the conscious awareness condition on epistemic states must be rejected, if we are to preserve the rationalising link between the epistemic and the practical.

How does all this relate to the catch-all propositions? As we saw earlier, rejecting the need for a catch-all proposition whilst preserving the plausible thought requires adopting a contextualist interpretation of the Bayesian formalism, which is associated with the concious awareness account of credences. But this account of credences stands in tension with a central assertion of Bayesian epistemology, namely that the relation between the epistemic and the practical is one of rationalisation, and as such, must be rejected. It follows that the contextualist interpretation too must be rejected. The usual interpretations by contrast are consistent with the claim that credences rationalise action; and on these interpretations, as I argued above, agents must be modelled with a catch-all proposition. This, of course, is not a decisive argument, but it does shift the burden of argument. And, as far as I am aware and with the exception of the contextualists discussed in this section, these interpretations are universally endorsed, and no criticism of them has been put forward that would threaten my argument for the inclusion of a catch-all proposition. Thus I conclude that the inclusion of such a proposition is mandatory in models of less-than-fully-aware agents. It follows that the contrivance claim is false: those who include catch-all propositions in their agents' algebras do not do so merely because it makes it possible to model awareness growth by refinement; rather such an inclusion is mandated by considerations about the nature of rationality.

4.

The mandatory inclusion of catch-all propositions is likely not to be welcome. Indeed, according to virtually everyone who has written on the topic, there is a serious problem with catch-all propositions: this is the deficiency claim. (Shimony, 1970; Glymour, 1980; Salmon, 1990; Earman, 1992; Henderson et al., 2010; Wenmackers and Romeijn, 2016; Bradley, 2017). The argument that all these writers take to establish this claim is, in the words of Bradley: "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?" (p. 255). In other words, it is uniquely difficult to know what credence to assign to catch-all propositions because, unlike other propositions that the agent considers, we do not know what they say. Our book-lover knows what *poems* says, but they do not know what *fiction that isn't novel or poems* says. So, how could they assign a particular credence to the latter? Thus the common argument for defectiveness is built on two claims: (1) there is a principled difference between regular and catch-all propositions, and (2) this difference entails a difference in tractability. In what follows, I argue against both of these claims.

Let us begin with the claim, (1) that there is a sharp distinction between regular catch-all

propositions. This supposed distinction is usually fleshed out in terms of ignorance about the content of catch-all propositions. For instance, Steele and Stefánsson write that "the agent has no idea how to specify the propositions' content" (forthcominga). But attempts to make this precise fall short. What constitutes knowledge of the content, that agents supposedly have about regular but not catch-all propositions? 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. Another 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 she was presented was a novel, but she would also presumably recognise whether a book she was presented with was neither a novel nor a poetry collection. Yet another option is that the agent understand the content of the proposition. This is not very precise, but in any case, it is too strong a requirement. 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 might have credences in *super-symmetry*, *quantum gravity*, and *other*. But, having no training in theoretical physics, I do not have any understanding of what *super-symmetry* or *quantum gravity* say; no more than I have an understanding of what *other* says. And I take it that, in this example, *super-symmetry* and *quantum gravity* are supposed to be actively considered, and *other* is supposed to be a catch-all proposition. If regular and catch-all propositions differ in how transparent their content is to the agent who considers them, it remains to be outlined precisely what this means.

Furthermore, considerations from the foundations of Bayesian epistemology should arouse suspicion that distinguishing between regular and catch-all propositions is infeasible. As we have already seen, the only requirement on entertained propositions in orthodox Bayesianism is that they form a Boolean algebra. Thus, there is nothing in the standard framework that could be drawn on to carve the distinction. Now of course, one might insist that the strength and stability of intuitions that there is a distinction along those lines is a reason to supplement the standard framework; rather than the characteristics of the framework being a reason to suppress these intuitions. However, these intuitions are, on closer inspection, much less stable than they are usually taken be. Take the book case, in which the entertained propositions are *novel*, *poems*, and *other*; and where, supposedly, the first two are regular and the third is a catch-all. The line of thought which I presume to have led to this characterisation of the case is the following: the agent enumerates the genres they know (the novel, then poetry), and then gets stuck—thus the rest of the epistemic space is filled with a proposition designated the catch-all. But it is easy to imagine a line of thought

with different intimations. Suppose for instance that the agent has no knowledge of experimental literature, but that their attention is directed towards that possibility—maybe because they harbour thrilling admiration for those whose interests are less dry than academic philosophy. When enumerating literary genres, this agent may very well start by listing the possibility that the book is "something uncommon and eccentric", and then list the banal options of *novel* and *poems*. Or, to take an even more basic example, consider again the umbrella case. The two propositions were described as *rain* and ¬rain. Presumably, ¬*rain* must be thought of as a catch-all: after all, it is what fills out the epistemic space against *rain*. But the umbrella case could just as easily have been described using the propositions *dry* and ¬*dry*; in which case, presumably, ¬*dry* would have been a catch-all proposition. What I hope this discussion shows is that it is much less clear than usually believed that the regular/catch-all distinction is cogent.[30]

But suppose that my objector does manage to make precise sense of the regular/catch-all distinction. He still needs to argue that (2) this distinction entails that considered propositions are more easily assigned a justified credence than catch-all propositions. I want to begin my discussion of this claim by presenting intuitions against it. Firstly, the book case: it is easy to imagine a situation where it is more difficult to determine how confident the agent ought to be in *poems* relative to novel *novels*, than it is to determine who confident they ought to be in *other* relative to *novel*∨*poems*. For instance, suppose that the agent has never spoken with their book-gifter friend about novels or poetry; they might have no idea how likely the two are relative to one another. Furthermore, suppose that the agent knows that their friend is an experimental literature enthusiast. Then, they might have good reason to believe that the book is likely not to be as banal as a mere novel or poetry collection. Secondly, the physics case: one might be justified (as indeed I take myself to be) in having a higher credence in the next fundamental physical theory than any of the current ones. By contrast, given my utter lack of knowledge of physics, I find determining the credence I ought to have in say *quantum gravity* as opposed to *super-symmetry* very difficult to determine. In light of these examples, we need to reassess the reliability of our initial intuition that the catch-all

[30] Steele and Stefánsson's rejection of catch-all proposition is based on a consideration different to the usual one I treat in this section. They write: "catch-all propositions are so abstract from the agent's point of view [...] that it is unclear [...] whether it is even cogent, to depict the agent as entertaining these propositions. [Indeed,] in order for an agent to make sense of a catch-all, she would presumably need to entertain some universal set of possibilities relative to which the catch-all can be defined as the complement of those possibilities she can properly articulate. But it is hard to see how the agent could have access to this universal set of possibilities (which might in fact not even be a coherent notion), given that, by assumption, some of these possibilities cannot be articulated." (They attribute this formulation of the argument to Alan Hájek.) The considerations above highlight flaws in this argument. Firstly, the requirement that agents be able to "articulate" the propositions they entertain is, I have argued, at best unclear. And secondly, their argument proves too much; for if they are right, the modelling of the umbrella case with the propositions *rain* and ¬*rain* is infelicitous. (I should note that Steele and Stefánsson's attitude to their argument is rather tentative—they themselves note related weaknesses.)

is supposedly intractable.

Let us now return to the fundamentals of Bayesian epistemology. When, according to Bayesians, are credences justified? There are many views on this question, the two main views being *subjectivism* and *objectivism*. According to subjectivists, an assignment of credences to an algebra of proposition is justified just in case it is probabilistically coherent. According to objectivists, an assignment of credences to an algebra of non-evidential propositions is justified just in case each of the atomic propositions receives an equal credence. Thus although thee proponents of these views are in many ways diametrically opposed, they are of one mind as far as the defectiveness claim is concerned: the difference between rational and irrational, or justified and unjustified, credences can be stated at the purely syntactic level; in other words, the content of the propositions and the agent's relationship to these contents is irrelevant to the justification of credences in them. For the subjectivist, what matters is that the credences in non-trivial propositions be additive, so the only criterion of rationality for our book receiver is that their credences in *novel*, *poems*, and *other* must sum to one. For the objectivist, what matters is that the credences in atomic be equal, so the only criterion of rationality for our book receiver is that their credences in *novel*, *poems*, and *other* must be equal. In other words, the credences that agents ought to have in various propositions is easily determined on both views, and the difference between catch-all propositions and actively considered ones does not arise.[31]

Where does this leave my opponent? In order to uphold the defectiveness claim, he must isolate the characteristic property of catch-all propositions, and propose an account of credal justification on which this characteristic is relevant to what agents ought to believe. This is a tall order.

5.

As I argued in §3, a less-than-fully-aware agent must always have a catch-all proposition in her algebra. Furthermore, as I argued in §4, there is nothing more wrong with catch-all propositions than with any other kind of proposition. So we can embrace the fact that our book lover's initial algebra must contain *poems*, *novels*, and *other*, and moreover, that upon becoming aware that the book could be of short stories, their subsequent algebra must then contain *poems*, *novels*, *stories*, and *other*. How is this related to the question of whether lateral awareness growth should be modelled by refinement or by expansion? As we saw in §2, is because they hold the conjunction of the

[31] This raises the question: whence the intuitions, both in favour of the defectiveness view (how am I to determine how likely it is that an unforeseen climate event takes place?) and against it (I am more confident in the next physical theory than in any of the current ones)? The best way to make sense of them, I suspect, is to recognise that they appeal to an implicit landscape of propositions that is much richer than what is explicitly mentioned. If this is indeed the case, then they are not after all cases of mere *lateral* awareness growth, and require a more complex treatment.

contrivance claim and the defectiveness claim that those who reject the refinement view do so. That is, they think that there is a problem with catch-all propositions (defectiveness), refuse to include catch-all propositions in an agent's algebra as a result, and therefore reject the refinement view (contrivance). So, as we reject the defectiveness and the contrivance claims, the reason to uphold the expansion view disappears. But it does not follow that the refinement view is true: the rejection of the contrivance and defectiveness claims does not entail the refutation of the expansion view.

In what follows, I present two related arguments for the refinement view, against the expansion view. Before doing so, a precise characterisation of these two views is required. The hallmark of the refinement view is that the later algebra $\mathcal{A}'$ is a fine-graining of the earlier algebra $\mathcal{A}$. This means that, for any proposition $A_i \in \mathcal{A}$, it must be that $A_i \in \mathcal{A}'$; and there are some propositions $A_j \in \mathcal{A}'$ such that $A_j \notin \mathcal{A}$. In more informal words, awareness growth on the refinement view consists in the splitting of one or more propositions; and lateral awareness growth specifically consists in the splitting of the catch-all proposition. So, on this view, what happens in the book case is that the agent's initial algebra is such that the atomic propositions are $\textit{novel}, \textit{poems}, \textit{other} \in \mathcal{A}$, and their algebra subsequent to becoming aware of *stories* is such that the atomic propositions are $\textit{novel}, \textit{poems}, \textit{stories}, \textit{other'} \in \mathcal{A}'$, where $\textit{stories} \vee \textit{other'} = \textit{other}$. Thus the meaning of the initial catch-all proposition can be captured as "something other than a novel or a poetry collection", and the meaning of the subsequent catch-all proposition can be captured as "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". By contrast to the refinement view, the expansion view is difficult to characterise positively. But in broad terms, the idea on this view is that the entirety of the underlying algebra is replaced with a different, richer algebra that contains the newly entertained proposition. Applied to the book case, this view yields that, if the initial algebra contains the atomic propositions $\textit{novel}, \textit{poems}, \textit{other} \in \mathcal{A}$, the later algebra contains the atomic propositions $\textit{novel'}, \textit{poems'}, \textit{stories'}, \textit{other'} \in \mathcal{A}'$.

This presentation of the views immediately suggests a first argument in favour of the refinement view, namely that it is only on this view that the agent can be said to retain an epistemic attitude towards the proposition they initially entertained as they undergo awareness growth. To see why, one must attend to the fact that propositions entertained by the agent are mandated to form a Boolean algebra. One of the consequences of this requirement is that propositions can be identified by their negations. So, in the book case, the catch-all proposition in the agent's initial algebra is identical to $\neg(\textit{poems} \vee \textit{novel})$. By contrast, the subsequent catch-all proposition is identical, not to $\neg(\textit{poems} \vee \textit{novel})$, but to $\neg(\textit{poems} \vee \textit{novel} \vee \textit{stories})$. Thus, the meaning of the catch-all proposition changes. This is captured on both the refinement and the expansion views: the agent

goes from entertaining *other* to *other'*. But, the difference between the views appears when we consider propositions besides the catch-all. Let us consider, for instance, the proposition that the book is either a novel or a poetry collection. This proposition remains is initially identified with $\neg$*other*, and subsequently, with $\neg$(*stories* $\vee$ *other'*). Now, the hallmark of the refinement view is the identification of *other* with *stories* $\vee$ *other'*. So, on this view, *poems* $\vee$ *novel* remains identical across awareness growth. But since the identification does not hold on the expansion view, the proposition *poems* $\vee$ *novel* does not remain identical—and the agent ceases to entertain the initial proposition, and begins entertaining a whole new one. Note that this argument can be run—albeit less tidily—with any proposition; thus what the argument really shows is that, on the expansion view, there is no proposition that the agent continues to consider across awareness growth. (I encourage readers who think visually to consult the diagram of §2 in which this argument is manifest.) Whereas, on the refinement view, every proposition in $\mathcal{A}$ persists in $\mathcal{A}'$; on the expansion view, no proposition in $\mathcal{A}$ persists in $\mathcal{A}'$.

And there is good reason to think that every instance of awareness growth does not constitute a complete revolution in the agent's awareness state. Indeed it is much more plausible that, rather than getting an entirely new map—an entirely different representation of the world—the agent who undergoes awareness growth comes to appreciate that their map was insufficiently precise in a place, and modifies this place as a result, retaining the rest of the map. So, on this rationale for the refinement view, awareness growth consists in coming to appreciate a distinction one was previously unable to make. Whereas at the earlier time, 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 discern a distinction among these genres, between short stories on the one hand, and yet-other genres on the other.

This case is made stronger by comparing lateral and internal forms of awareness growth—this is my second argument for the refinement view. In §2, I introduced what I called "internal" awareness growth, which I illustrated with the following modified book case. Consider an agent who initially entertains the propositions that the book of fiction might be a novel or a book of poems, and subsequently realises that, if it is a book of poems, it might be a book of prose poems or a book of poems in verse. (We can falsely assume that this is an exhaustive disjunction for simplicity.) As I asserted in §2, everyone agrees that this kind of awareness growth should be modelled by refinement: the agent's initial algebra $\mathcal{A}$ has for atomic propositions *poems*, *novel*, and *other*; and their subsequent algebra $\mathcal{A}'$ also has *novel* and *other*, but instead of *poems*, it has *prose* and *verse*. (The agent continues to consider the proposition *poems*, not as an atomic proposition, but as the disjunction of *prose* and *verse*.) The rationale for modelling the agent in this way is obvious: the agent becomes able to differentiate between two types of poetry, and as such, modifies their map

in that place to reflect this newfound awareness. This is the same rationale that I gave above for modelling lateral awareness growth by refinement—there too, the agent's powers of discernment increases. But remember that I dubbed my distinction between different forms of awareness growth "pre-theretical". Later, in §4, I argued that there is no principled way to distinguish between catch-all propositions and "regular" ones. It follows that what appeared, pre-theretically, to be a distinction—namely, the lateral/internal distinction—cannot in fact be drawn. And given, as is uncontroversial, that internal awareness growth should be modelled by refinement, so too must lateral awareness growth. Again: given that there is no normatively relevant difference between "regular" propositions and catch-all propositions, and given that appealing to such a difference would be the only way to create a rift between internal and lateral forms of awareness growth, we must conclude that the two ought to be modelled identically: by refinement.

6.

To conclude: we must accommodate awareness growth within Bayesian epistemology, and we must do so by refinement. As is widely recognised, this requires the inclusion of a catch-all proposition; but contrary to widespread belief, such an inclusion is required anyway, and does not pose a specific problem.

9th November 2020

References

Bradley, R. (2017). *Decision Theory with a Human Face.* Cambridge University Press.

Bronfman, A. (2014). Conditionalization and Not Knowing That One Knows. *Erkenntnis 79*(4), 871–892.

Buchak, L. (2013). *Risk and Rationality.* Oxford University Press.

Carr, J. (2015). Epistemic Expansions. *Res Philosophica 92*(2), 217–236.

Chalmers, D. J. (2011). Frege's Puzzle and the Objects of Credence. *Mind 120*(479), 587–635.

de Finetti, B. (1974/1990). *Theory of Probability.* Wiley.

Dietrich, F. and C. List (2016). Mentalism Versus Behaviourism in Economics: A Philosophy-of-Science Perspective. *Economics and Philosophy 32*(2), 249–281.

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.

Eriksson, L. and A. Hájek (2007). What Are Degrees of Belief? *Studia Logica 86*(2), 185–215.

Frege, G. (1892). On Sense and Reference. In P. Geach and M. Black (Eds.), *Translations from the Philosophical Writings of Gottlob Frege*. Oxford University Press. 1980.

Frigg, R. and S. Hartmann (2020). Models in Science. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Gallow, J. D. (forthcoming). Updating for Externalists. *Noûs*.

Glymour, C. (1980). *Theory and Evidence*. Princeton University Press.

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.

Hájek, A. and H. Lin (2017). A Tale of Two Epistemologies? *Res Philosophica 94*(2), 207–232.

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.

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.

Maher, P. (1995). Probabilities for New Theories. *Philosophical Studies 77*(1), 103–115.

Mahtani, A. (2017). The *Ex Ante* Pareto Principle. *Journal of Philosophy 114*(6), 303–323.

Mahtani, A. (forthcominga). Awareness Growth and Dispositional Attitudes. *Synthese*.

Mahtani, A. (forthcomingb). Dutch Book and Accuracy Theorems. *Proceedings of the Aristotelian Society*.

Okasha, S. (2016). On the Interpretation of Decision Theory. *Economics and Philosophy 32*(3), 409–433.

Pettigrew, R. (2016a). *Accuracy and the Laws of Credence*. Oxford University Press.

Pettigrew, R. (2016b). Accuracy, Risk, and the Principle of Indifference. *Philosophy and Phenomenological Research 92*(1), 35–59.

Pettigrew, R. (forthcoming). Logical Ignorance and Logical Learning. *Synthese*.

Ramsey, F. P. (1926/1931). Truth and Probability. In *The Foundations of Mathematics*. Kegan Paul.

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.

Schoenfield, M. (2017). Conditionalization Does Not (in General) Maximize Expected Accuracy. *Mind 126*(504), 1155–1187.

Shimony, A. (1970). Scientific Inference. In R. Colodny (Ed.), *The Nature and Function of Scientific Theories*, pp. 79–172. University of Pittsburgh Press.

Steele, K. and H. O. Stefánsson (forthcominga). Belief Revision for Growing Awareness. *Mind*.

Steele, K. and H. O. Stefánsson (forthcomingb). *Beyond Uncertainty*. Cambridge University Press.

Tang, W. H. (2012). Regularity Reformulated. *Episteme 9*(4), 329–343.

Titelbaum, M. G. (2012). *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford University Press.

Wenmackers, S. and J.-W. Romeijn (2016). New Theory About Old Evidence. *Synthese 193*(4), 1225–1250.