Awareness revision and belief extension

Joe Roussos

Institute for Futures Studies

This is a pre-copy-editing version of a paper which is forthcoming in the *Australasian Journal of Philosophy.* Please do not distribute and cite the published version.

Abstract

What norm governs how an agent should change their beliefs when they encounter a completely new possibility? Orthodox Bayesianism has no answer, as it takes all learning to involve updating prior beliefs. A partial proposal is Reverse Bayesianism, which mandates the preservation of ratios of prior probabilities, but it faces counterexamples introduced by Mahtani (2021). I propose to separate awareness growth into two stages: awareness revision and belief extension. I argue that Mahtani's cases highlight that we need to theorize awareness revision before we can define a proposal for belief extension, such as Reverse Bayesianism. I provide a formal model of awareness revision which makes explicit how propositions are distinguished within awareness states and identified across them. Reformulating Reverse Bayesianism to take input from my model allows it to navigate Mahtani-style cases. My model leaves open how agents choose to identify propositions across awareness states, and I propose that they ought to do so conservatively: preserving undisturbed prior reasoning about the structure of their awareness. I then spell out this proposal in a special case. This is a partial proposal, and I close with a discussion of how to elaborate on it and how to advance research into awareness revision.

1. Introduction

Ordinary people like you and me regularly confront new possibilities. When I moved to Stockholm, I learned that Riddarholmen lies just west of Gamla Stan. Prior to this I had never heard of either of these places, and it is natural to say that I had no attitude whatsoever to propositions regarding their relative positions. I was simply *unaware* of these propositions and their objects. Upon coming to Stockholm, my awareness grew to encompass various possibilities regarding the positions of Gamla Stan and the surrounding islands, and I formed my new belief.

How does rationality govern this kind of awareness growth? It is a strange failing of our most popular formal models of belief that they have little to say about it. I speak of Bayesian models, which represent beliefs with probabilities and learning with conditioning. Bayesianism is a rich and successful theory, but in Bayesian models all resolutions of uncertainty take place by updating pre-existing beliefs. Agents must have priors for propositions to learn about them at later stages and so Bayesianism leaves no room for agents to learn about genuinely new states of affairs, and has no guidance for real agents when they undergo such changes of awareness. New possibilities are at once so common and so bound up in our most pressing epistemic challenges that filling this gap is a matter of first importance to epistemology.

In this paper I aim to advance the study of awareness growth in the context of probabilistic beliefs in three ways. First, I conceptually separate the way an agent accommodates awareness growth, which I call awareness revision, from how they use their prior beliefs to reason about their new set of possibilities, which I call belief extension. These two processes are distinguished from the well studied process of belief revision (for example, by Bayesian conditioning). Second, I use this distinction to argue that a currently popular approach to awareness growth, called Reverse Bayesianism, is a model of belief extension which neglects important questions about awareness revision. This clarifies a recent debate between Bradley (2017, 2022) and Mahtani (2021), and resolves the problem cases Mahtani introduces. However, an important new question is raised: when an agent's awareness grows, how do they relate the possibilities they were previously aware of to those in their new possibility space? My third contribution is to sketch what an answer to this question might look like, in the

specific case of an agent who makes judgements of equipossibility and uses the principle of indifference to assign probabilities to those possibilities.

A quick preview of the new concepts I will use: *Awareness revision* is the formation of a new subjective awareness state, or space of subjective possibilities. It involves reasoning about the content of and logical relations between propositions, without taking an attitude to them more committal than awareness qua possibilities. *Belief extension* is the formation of a provisional credal attitude on the new space of possibilities. It involves determining what one's old probabilistic beliefs have to say about the new awareness state. Belief extension is distinguished from *belief revision* in that it involves no evidence about what is the case: an agent who learns that a new proposition is possible, without learning anything about the likelihood of that or other propositions, undergoes belief extension but not belief revision. Realistic cases of awareness change typically involve all three processes—awareness revision, belief extension, and belief revision—but the latter is not my topic here.

In the rest of this introduction I situate my project in the growing literature on awareness, and then outline the plan of the paper. As my opening example indicates, unawareness refers to a 'lack of conception rather than a lack of information' (Schipper 2014: 1). Representing these two kinds of uncertainty in a one model is an important topic in information economics and computer science. A central result by Dekel, Lipman, and Rustichini (1998) shows that a model with a single standard state space (that is, one in which events are sets of worlds) cannot represent unawareness.[1] One can do better with a syntactic model (Fagin and Halpern 1987), or by introducing multiple state spaces (Heifetz, Meier, and Schipper 2006). In economics especially the interest is in models which can represent multiple agents with different states of awareness; for example, for the purpose of studying speculative trade on the basis of asymmetric information (Heifetz, Meier, and Schipper 2013). Schipper (2015) provides a review of these literatures, noting the important links with modal and epistemic logic.

[1] Representing unawareness here means having an explicit unawareness operator which, together with a knowledge operator, allows for the expression in the model of events like "the agent is unaware of E" and "the agent knows that they are unaware of E".

My topic is instead the *changes* of awareness experienced by a single agent. Like Heifetz, Meier, and Schipper (2006), I will make use of multiple state spaces, although in my case they represent different states of awareness for a single agent. Hill (2010) studies single-agent awareness growth in a logical setting, developing a model of and logic for awareness change. I am inspired by several parts of Hill's treatment: the adoption of the agent's limited perspective and the separation of the awareness and belief changes the agent undergoes. My setting is different in that I study models of probabilistic belief, and so the belief revision with which my model interfaces is quantitative rather than qualitative. I focus on awareness and belief, and do not discuss preference or choice. Nonetheless, I spend some time discussing a proposal initially developed by the economists Karni and Vierø (2013), who studied awareness growth in the context of choice. Their proposal, called "Reverse Bayesianism", preserves the ratios of probabilities of propositions which the agent was previously aware of. Philosophers have been inspired by this aspect of their model, and discussed it separately from the complex choice-theoretic underpinnings of Karni and Vierø (Bradley 2017; Mahtani 2021; Steele and Stefánsson 2021b, 2021a). Bradley's presentation of Reverse Bayesianism, discussed below, has some similarities with an earlier proposal due to Williamson (2003), who framed the problem as one of language change. The latter is part of a tradition of discussing this limitation of Bayesianism in terms of language and interpretation shifts—including Lakatos (1968) on Carnap's confirmation function and language change, and more recently Romeijn 2012.

Mahtani, Steele and Stefánsson frame their discussion as concerning the updating of credences, and in doing so they are naturally in conversation with philosophers of science who study how confirmation theory should handle new hypotheses (for example Earman 1992; Romeijn 2005; Wenmackers and Romeijn 2016). There are links here to the literature framed in terms of language change—for example, Gillies (2001) discusses Bayesian statistics's dependence on a fixed set of hypotheses with reference to Lakatos 1968. A significant proposal for handling awareness growth in a Bayesian setting comes from Shimony (1970), who proposed that a rational scientist should reserve some probability for a catch-all hypothesis representing an as yet unknown theory. This allows him to model awareness growth as the specification of this catch-all (or a part of it), which accompanies belief revision via standard Bayesian update of the catch-all prior. Criticism of this approach highlights the opaque nature of the catch-all,

and its collapse of all unawareness into conscious unawareness (Steele and Stefánsson 2021b).[2] I do not use any prior representation of potential new possibilities. Awareness change is modelled using two awareness states, meaning that the agent need have no prior awareness of their unawareness.

In contrast with this focus on belief revision, one of my contributions is to illustrate the importance of studying the change of awareness itself. I show that Reverse Bayesianism is a theory of belief extension, rather than belief revision as Steele and Stefánsson (2021a) take it to be. Indeed, I barely discuss belief revision (that is, Bayesian conditioning or a generalisation), noting only how my model of awareness revision and my adjusted version of Reverse Bayesianism might plug into various approaches to updating beliefs following awareness growth.

Two final comments on the literature and my methodology. First, I follow Hill (2010) and Bradley (2017) in adopting a first-personal perspective, modelling situations of awareness growth as the agent experiences them. Thus, my models have no more information in them at any stage than the agent has. I prefer this to more omniscient third-personal models, because I am interested in generating norms for awareness growth that are usable by real agents who undergo it. If one works from an omniscient perspective, it is easy to lose track of what is interesting and challenging about awareness growth. In taking awareness growth seriously and modelling it from the agent's perspective, my model is also fully subjective. The possibilities in the model are subjective, and the tautological and contradictory proposition are merely those that the agent takes to be tautological or contradictory, given their state of awareness. Steele and Stefánsson (2021a, 2021b) also make use of subjective possibilities, as does some work in computer science and logic (cf. Schipper 2015). The philosophy of science literature (for example, Wenmackers and Romeijn 2016), by contrast, tends to use objective possibilities, while the economics literature (for example, Heifetz, Meier, and

[2] Conscious unawareness is the state in which an agent is aware than there is something they are unaware of. It is distinguished from the presumably more common state of unconscious unawareness, in which the agent is unaware of their unawareness.

Schipper 2006) naturally adopts a third-personal perspective suitable to the project of studying multiple agents.

Second, I use a novel formalism for modelling awareness and belief. In place of the standard framework in which propositions are sets of possible worlds, I use a model in which propositions are basic. Probabilities are defined on lattices of propositions. This allows me to track the changes in logical understanding that the agent undergoes as part of their awareness growth. In standard set-based models, by contrast, the logic is built into the set operations and so it is difficult to hold on to the same propositions while altering their logical relations.

Here is the plan for the rest of the paper. In §2, I introduce awareness revision and belief extension. §3 introduces Reverse Bayesianism and examines Mahtani's challenge to it. §4 develops my approach to awareness revision. In §5, I briefly show how this approach resolves Mahtani's problem case, before turning in §6 to a discussion of the core problem of awareness revision: how to identify propositions across states of awareness. §7 concludes.

2. Awareness growth

My focus in this paper is on one species of awareness change, awareness growth, in which agents become aware of new possibilities. (The complementary form of awareness change, in which agents lose track of possibilities or remove them from consideration, will not be discussed.) Let us begin by noting that there are several ways that one's awareness can grow. These have been characterized with reference to partitions of the space of possibilities an agent is aware of. Propositions in a partition are mutually exclusive (the conjunction of any two is contradictory) and collectively exhaustive (the disjunction of all propositions in the partition is tautological).

Imagine that Naledi is a South African who has just moved to Sweden. She is considering tomorrow's weather and, being South African, she is aware of two possible kinds of weather: rain and sun. Let us use RAIN to represent the proposition that it will rain tomorrow, and likewise for SUN. For simplicity's sake, we will suppose that for her {RAIN, SUN} is a partition. We will imagine that Naledi, unlike real South Africans, has never even heard of the various forms of terrible northern European weather that she is about to experience.

We can distinguish three ways in which Naledi's awareness might grow. First, Naledi could become aware of snow, taking it to be a third kind of weather, and thus expanding the partition of weather states from {RAIN, SUN} to {RAIN, SUN, SNOW}. Second, she could realize that RAIN and SUN are not mutually exclusive after all, perhaps by observing a sun shower.[3] She now distinguishes ¬RAIN from SUN, so that {RAIN, ¬RAIN} and {SUN, ¬SUN} are independent partitions. Third, she could realize that she needs to worry about temperature, since it can be either hot or cold with either rain or sun. The propositions HOT and COLD form a partition, and the finest partition she recognizes is now {RAIN∧HOT, RAIN∧COLD, SUN∧HOT, SUN∧COLD}.

The first kind of awareness growth is called *expansion*, and it involves recognising that the initial set was not a partition because it was not exhaustive. This requires altering several logical relations: the disjunction RAIN∨SUN is no longer equivalent to ⊤, the tautology, and the negation ¬RAIN is no longer equivalent to SUN.[4] The second kind of awareness growth described above does not seem to have a common name and indeed has been relatively neglected. I will call it *clarification*. Clarification also involves recognising that the initial set was not a partition, but here because the propositions were not pairwise disjoint.[5] So the conjunction RAIN∧SUN is no longer equivalent to ⊥, the contradiction and again the negation ¬RAIN is no longer equivalent to SUN. The third kind of awareness growth is called *refinement*, and it is distinct in that it does not involve realising that a set is not a partition. Instead, refinement involves the

---

[3] Which South Africans delightfully refer to as a "monkey's wedding".

[4] A set of propositions $\mathbb{X}$ is *exhaustive* when the disjunction of all its elements is the tautology: $\bigvee\mathbb{X} = \top$. Expansion occurs when the agent realizes this condition fails and introduces a new proposition to create a partition.

[5] A set of propositions $\mathbb{X}$ is *pairwise disjoint* when, for any two elements X and Y of $\mathbb{X}$, $X \wedge Y = \bot$. Clarification occurs when the agent realizes that this condition fails and forms new partitions accordingly. This can involve realising that one or more pairs of propositions have non-contradictory conjunctions.

introduction of a new partition which is finer than an existing partition.[6] So, for example, RAIN∧HOT implies RAIN, and {RAIN, SUN}—while still a partition—is no longer the finest partition that Naledi recognizes.

A brief comment about clarification, which I have not seen discussed before. Clarification might look like it is reducible to expansion, refinement, or a combination of the two. A reader might be inclined to reframe Naledi's prior awareness as involving the possibilities RAIN∧ ¬SUN and ¬RAIN∧SUN, and so view the experience as an expansion. Or one might note that, after the clarification, Naledi's awareness looks refined, in that she now has a partition {RAIN∧SUN, RAIN∧ ¬SUN, ¬RAIN∧SUN, ¬RAIN∧ ¬SUN}. This misses something important about clarification, however. Expansion involves changes to one's conceptual understanding, such that one better understands how they relate to the world and in particular to what is necessarily true of it—that is, of how certain concepts relate to the tautology. By contrast, clarification involves changes to one's conceptual understanding such that one better understands what is necessarily false of the world—that is, of how certain concepts relate to the contradiction. Both expansion and clarification can be analysed as refinements of a sort, by imputing various attitudes to the agent. But this denies the quite common experience of changes in conceptual understanding which lead to revisions of logical relations between certain concepts. I prefer a first-personal picture on which we recognize that for Naledi, RAIN and SUN were genuinely exclusive and exhaustive options before, and that the way in which her awareness changed is connected to the realisation that RAIN∧SUN is not contradictory. These conceptual differences from expansion and refinement (considered individually) also count against reducing clarification to a combination of the two.

When someone like Naledi undergoes awareness growth they come to hold beliefs about their new possibilities.[7] I would guess that in most real cases, agents undergo

---

[6] A partition $\mathbb{X}$ is *finer than* a partition $\mathbb{Y}$ when, for every proposition Y in $\mathbb{Y}$, there is a proposition X in $\mathbb{X}$ such that $X \vDash Y$. Refinement is also sometimes called "specification".

[7] I do not address full belief in this paper, and so wherever "belief" occurs below it refers to credence, also known as degree of belief.

changes of awareness because they gain some evidence about the new possibilities, which then informs the beliefs they form. For example, Naledi might see a weather report which assigns a 30% chance to SNOW.[8] In so doing she does not just become aware of the possibility of snow, she also gains evidence for its likelihood. To her, the experience of becoming aware of SNOW will be bound up with the experience of coming to hold beliefs about SNOW and revising her beliefs about RAIN and SUN. However, for the purpose of analytical clarity I want to separate out different processes which occur in such cases.

I propose to analyse awareness change experiences as consisting of three stages. The first stage is a kind of reasoning about the new possibilities which takes place without any attitude more committal than awareness. This stage, which I call awareness revision, is the main focus of this paper. Awareness revision is constituted by the introduction of new propositions, the introduction of new logical relations between these propositions and the old propositions, and (sometimes) changes to the agent's understanding of the logical relations between old propositions. There are different kinds of awareness revision, corresponding to the kind of experience the agent has. Above I introduced three such kinds. To clarify how the 'experience' and 'revision' terminology works, here is an example for expansion: we will say that 'the agent has an expansion experience', naming a kind of experience in which awareness changes, and that 'their awareness undergoes expansion'—the latter use of 'expansion' naming a type of awareness revision.

Once Naledi understands the logical relations that hold between her new possibilities, she can reason about what credal attitude to take towards them. Naledi had credences about RAIN and SUN, and presumably those attitudes were based on some evidence which may still be relevant. Then there is whatever new evidence she has gained in the experience that made her aware of SNOW. This needs to be incorporated into her new credences. I will separate these different belief-related aspects of awareness change. In stage two, which I call belief extension, the agent works out what their prior credal

[8] For a discussion of cases like this, which involve both deference to an expert's probabilities and awareness growth, see Roussos 2021.

attitudes imply about their new state of awareness. The result of this process is a kind of provisional attitude, an extended prior belief state. In stage three the agent updates this provisional attitude with whatever new evidence was gained during the experience—for example, the weather app's reported chance of snow. As is standard, I refer to this process of updating in light of new evidence as belief revision.

I propose this division into three stages for the following reasons. Awareness revision requires a kind of logical reasoning that is not part of standard Bayesian models of belief. Those models take agents to be logically omniscient and assume that all reasoning takes place with reference to a fixed background logic which captures the actual logical relations between propositions. Second, it is possible, though perhaps uncommon, to become aware of a possibility without learning anything about its likelihood. Such cases can then be modelled as involving belief extension but not belief revision. On the other hand, in cases where the agent does learn information about the likelihood of the new possibilities, this separation allows us to employ our familiar thinking about updating a prior set of beliefs even though, strictly speaking, there were no beliefs about the new possibilities prior to the experience. This is because we represent the agent as first extending their prior beliefs to cover the new set of possibilities and then revising these provisional attitudes.

3. Reverse Bayesianism

My goal is to provide a model for awareness revision, but I will set up the motivation for it with reference to a popular proposal for belief extension, called Reverse Bayesianism. This is because I am inspired by a recent paper by Mahtani (2021), which presents a challenge to Bradley's (2017) version of Reverse Bayesianism.

Once an agent's awareness grows, they come to have beliefs about the possibilities that they now entertain. I will assume a broadly Bayesian picture on which agents start and end with probabilistic beliefs. So, the question of rational belief extension is: what constraint does rationality put on the agent's assignment of credence to new possibilities, in light of their prior credences? So, in my example: what can we say about how Naledi should assign credence to SNOW, after expansion; or to RAIN∧SUN, after clarification; or to HOT and COLD, after refinement? Reverse Bayesianism provides an

indirect answer to this question. Here is the definition of Reverse Bayesianism that I will use.

> **Reverse Bayesianism:** When an agent undergoes awareness growth, their belief state should be rigidly extended to the new possibility space.

The core notion is *rigid extension*, which comes from Bradley (2017: 258) and is defined as follows. Let $\mathcal{X}$ be the set of all propositions of which the agent is initially aware, so that $\bigvee\mathcal{X}$ is the tautology (according to the agent). Let P be the probability function on $\mathcal{X}$ representing their credences. The agent becomes aware of some new propositions contained in the set $\mathcal{E}$. We form a new set $\mathcal{Y}$, the closure of $\mathcal{X} \cup \mathcal{E}$ under the Boolean operations. Note that $\bigvee\mathcal{X}$ is in $\mathcal{Y}$. We now consider all probabilities defined on $\mathcal{Y}$.

> **Rigid Extension:** For any P, a probability function P^+ defined on $\mathcal{Y}$ is called a *rigid extension* of P to $\mathcal{Y}$ iff, for all $X \in \mathcal{X}$, $P^+(X|\bigvee\mathcal{X}) = P(X)$.

So, for any proposition X that you were previously aware of, your conditional probabilities for that proposition, conditional on the proposition you previously took to be tautological, should equal your prior unconditional probability in X. For example, if $\mathcal{X} = \{\text{RAIN,SUN}\}$ is Naledi's initial awareness state[9], then Reverse Bayesianism says that $P^+(\text{RAIN}|\bigvee\mathcal{X}) = P(\text{RAIN})$, and the same for SUN. In expansion cases, $\bigvee\mathcal{X}$ is not equivalent to the tautology after the awareness change, and so the unconditional probabilities of known propositions, such as $P^+(\text{RAIN})$, can change. (The same happens in clarification cases, though these have not been discussed heretofore in the literature.) In refinement cases, where $\bigvee\mathcal{X}$ is still the tautology, the credences don't change at all for known propositions, and the only restriction on assigning credence to HOT and COLD is rigid extension.

Reverse Bayesianism has a signature feature: the requirement that the agent preserves the ratios of probabilities of propositions they were previously aware of. If Naledi

[9] An awareness state must be closed under the Boolean operations, so strictly speaking Naledi's prior awareness state is $\text{cl}(\{\text{RAIN,SUN}\}) = \{\bot, \text{RAIN,SUN}, \top\}$. For brevity I will refer to the relevant partition as the awareness state when there is no ambiguity.

previously assigned $P(\text{SUN}) = P(\text{RAIN}) = 0.5$, then this 1:1 ratio between RAIN and SUN should be preserved. [10]

Defenders of Reverse Bayesianism justify it on grounds of conservativity of belief change, which underlies norms for rational belief revision.[11] In the face of new information, a rational agent updates their beliefs to accommodate what has been learned but does no more; going beyond the demands of the evidence may lead to unnecessary information loss and is unjustified. Bradley considers cases of expansion and refinement, and in each case he argues that the form of belief change that the agent undergoes is not one which requires any change in their comparative attitudes between known propositions. So, in my example, when Naledi's awareness changes, she learns nothing to imply that her prior judgement about the relative plausibility of rainy vs. sunny weather should be revised (Bradley 2017: 255). Rigid extension guarantees that all prior relational beliefs are preserved. I will return to this motivation in §6.

3.1. Mahtani's challenge to Reverse Bayesianism

In a recent paper, Mahtani (2021) presents a problem for Reverse Bayesianism based on cases in which it seems to disallow assigning positive credence to new possibilities.

---

[10] Some take this ratio-preservation requirement to be definitive of Reverse Bayesianism. As noted in §1, Reverse Bayesianism was developed by Karni and Vierø (2013) in a highly constrained setting focused on choice, and it is there that the ratio rule originated. Related approaches have been advanced in philosophy by Wenmackers and Romeijn (2016) and Bradley (2017), and applied by Vallinder (2018) and Roussos (2021). I develop Bradley's version and will simply refer to the approach I outline above as "Reverse Bayesianism", without trying to track the differences and disputes amongst Reverse Bayesians.

[11] Vallinder (2018) and Steele and Stefánsson (2021a, 2021b) also characterize Reverse Bayesianism this way. Another notable justification is Karni and Vierø's axiomatisation of Reverse Bayesianism, but theirs is a complex, choice-theoretic presentation and applies only to certain kinds of awareness change. It would take us too far afield. Suffice to say that it involves strong assumptions, including that the agent has a precise probability and utility function after the awareness change. Agents which meet these and other constraints are required to obey the ratio rule (Karni and Vierø 2013).

These cases turn on propositions which appear to split during the growth of awareness, so that, looked at one way, the case involves expansion, while looked at another way, it involves refinement. Here is the first case.

> *The Other Tenant*. You are staying with Bob, who lives with his landlord. You hear someone singing in the shower and wonder who it is. You find two propositions equally likely: that the singer is Bob and that the singer is the landlord. Since you know that Bob is a tenant, you also have credence in the proposition that the singer is a tenant. Now suppose that it occurs to you that there might be another tenant living in the same flat, and that they might be the singer. (Paraphrased from Mahtani 2021: 8984–85)

Let us use the labels BOB, LANDLORD and TENANT respectively for the propositions that the singer is Bob, the landlord, and a tenant. Initially, you assign BOB and LANDLORD probability 0.5. You also assign 0.5 to TENANT, since it is equivalent to BOB given your initial state of awareness. Then, you become aware of the proposition that the singer might be another tenant, which I will label OTHER. Once you do, how should your credence change?

Reverse Bayesianism says that you should preserve the ratio of probabilities of known propositions. So, the LANDLORD:BOB ratio of 1:1 should remain fixed. This seems intuitively sensible. The trouble comes from TENANT and BOB. Since TENANT and BOB are propositions you were aware of previously, Reverse Bayesianism presumably applies to them. This means that the TENANT:BOB ratio ought to remain 1:1. But after your awareness grows, you recognize that OTHER entails TENANT and that it is disjoint from BOB. Presumably, in so doing, you assign some credence to OTHER as it is a contingent possibility for you. But if the TENANT:BOB ratio must remain 1:1, OTHER can have no credence assigned to it at all. Reverse Bayesianism is doing too much: forbidding the assignment of positive credence to a new possibility.

There are a few interesting aspects of this case. It puts pressure on the notion of a 'familiar proposition'—TENANT is in one sense a familiar proposition, as the agent (you) did previously recognize the possibility that the singer was a tenant. But its contents have changed to encompass the possibility that this tenant is not Bob. It also illustrates that whether an awareness change is an expansion or a refinement depends

on how we label propositions: 'we have a refinement relative to the possibilities LANDLORD and TENANT, but an expansion relative to the possibilities LANDLORD and BOB' (Mahtani 2021: 8986). This is puzzling since, as I pointed out above, these two kinds of awareness change involve quite different changes of conceptual understanding and require different revisions of one's understanding of the logical relations between propositions.

I take Mahtani's challenge to be a serious one, and I agree that Reverse Bayesianism does not give a sensible recommendation here. But I see the trouble as coming from the lack of clarity about which propositions are identical to which others, exemplified by the seeming freedom to characterize the case as a refinement or expansion. This is a matter of awareness revision, which comes prior to belief extension. I will provide a model of awareness revision which resolves them, and thus shows that the issue was not with Reverse Bayesianism per se—indeed, suitably adjusted Reverse Bayesianism can handle these cases fine when applied in conjunction with my awareness revision model. It is not all good news for Reverse Bayesianism though: exploring the nature of awareness revision will delimit a set of circumstances under which it is a good rule to follow.

4. Awareness revision

I will use a model in which propositions are primitive objects, rather than sets of possible worlds as is standard. Propositions are collected in a set $\mathcal{A}$, which is endowed with two binary operations: $\wedge$, called meet, and $\vee$, called join.[12] Together, the set and the two operations form an algebraic structure $\mathbb{A} = (\mathcal{A},\wedge,\vee)$ called a lattice. One important kind of lattice is a complemented distributive lattice, also known as a Boolean algebra.[13]

---

[12] Meet and join are associative and commutative, each is idempotent, and they obey an absorption law. Idempotency means that $X \vee X = X$, and $X \wedge X = X$. The absorption laws are $X \vee (X \wedge Y) = X$, $X \wedge (X \vee Y) = X$.

[13] In a *distributive* lattice, the meet and join operations distribute over one another. A *bounded* lattice has two special elements, denoted $\bot$ and $\top$ and called bottom and top respectively; defined by $X \wedge \bot = \bot$, and $X \vee \top = \top$ for all $X \in \mathcal{A}$. A *complemented* distributive lattice is a

Complementation, ¬, is a unary operation and we can denote a Boolean algebra as $\mathbb{A} = (\mathcal{A}, \wedge, \vee, \neg)$. Boolean algebras are important because one can define probabilities on them, and thus use them for modelling beliefs. They are called *Boolean* algebras because they correspond to propositional logic: we interpret meet and join as the logical operations of conjunction and disjunction, and complementation as negation. So $(\mathbb{A}, P)$ represents the agent's belief state and in particular $\mathbb{A}$ represents their state of awareness, including the logical relations they take to hold between propositions.

Usually, this kind of model is equivalent to a 'sets of possible worlds' model: Stone's representation theorem shows that every Boolean algebra is isomorphic to a field of sets.[14] I make the distinction because, as I showed above, changes of awareness involve re-evaluating the logical connections between propositions. My model of awareness revision involves subjectively identifying propositions across states of awareness and preserving some but not all of their relations. For example, as we saw above, expansion preserves conjunction and disjunction across awareness states but not negation. We can represent this using different lattices, with distinct meet, join, and complementation operations, and then analyse the relations between such structures using tools from abstract algebra. Set theory, on the other hand, typically has fixed operations of union, intersection, and complementation. While it is possible to represent changes of awareness in such models, it is much harder to represent them *as* changes in logical understanding.

Two points on the formalism before we get to the model: as I will be considering multiple Boolean algebras, the ⊤ and ⊥ notation is ambiguous. I will therefore denote the top and bottom of $\mathbb{A}$ as $\bigvee\mathcal{A}$ and $\bigwedge\mathcal{A}$ respectively. I will also sometimes talk in terms of an implication relation, ⊨, defined by $X \models Y$ iff $X \wedge Y = X$ iff $X \vee Y = Y$. ⊨ is also called the *order* for the lattice.[15]

---

bounded lattice such that, for each $X \in \mathcal{A}$, there is a unique element of $\mathcal{A}$, denoted $\neg X$, such that $X \wedge \neg X = \bot$ and $X \vee \neg X = \top$.

[14] And, of course, fields of sets are Boolean algebras in the algebraic sense.

[15] Some may be more familiar with the order-theoretic perspective (with which there is no conflict). I prefer to highlight the algebraic properties of lattices because it makes the link with

I will make use of a visualisation technique for lattices, called a Hasse diagram. A Hasse diagram has a node for each element of the lattice and draws a line that goes upward from X to Y whenever $X \vDash Y$. All that matters is the start- and end-points of a line, there is no interpretation of lines crossing. Two simple lattices are shown in Figure 1. Note one differences between them: in Figure 1(a), $P \vee Q = \top$, so if we take this to be a *complemented* lattice, $\neg P = Q$. In Figure 1(b), by contrast, $P \vee Q \neq \top$, and instead $\neg P = Q \vee R$.

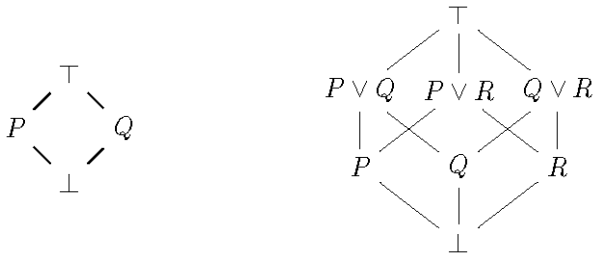

Figure 1. Hasse diagrams showing two simple lattices. (a) Two non-trivial elements in a single, two-element partition. (b) Three atoms in a single, three-element partition.

The initial and final awareness states will be represented by Boolean algebras $\mathbb{A} = (\mathcal{A},\wedge,\vee,\neg)$ and $\mathbb{B} = (\mathcal{B},\sqcap,\sqcup,\sim)$. Note that I have used different notation for the operations in $\mathbb{A}$ and $\mathbb{B}$: $\sqcap$, $\sqcup$ and $\sim$ are the meet, join and complement of $\mathbb{B}$; they are defined as above but relate elements of $\mathcal{B}$ rather than $\mathcal{A}$. I use this notation to emphasize that these are introduced as distinct structures, which we look for a relation between. The problem of awareness revision is to find a relation between the old and new algebra which represents the agent's judgement of which propositions in the new algebra correspond to which propositions in the old algebra. As a modelling task, getting this right involves matching our intuitions about what is changed (such as complementation relations) and what remains the same (where this includes both propositions and logical relations between them), as discussed above.

For this we need a mapping between the two algebras. As we're considering awareness growth, where new possibilities are added, we need an injective map. As logical

---

propositional logic clearer, and I find it fits more naturally with concerns about morphisms, to be introduced below.

relations are represented by the algebraic structure of the lattice, and as some of this is preserved, it is natural to look at the class of structure-preserving maps called lattice homomorphisms. As we start and end with a Boolean algebra, one might think to use a Boolean algebra homomorphism. But this would be a mistake, since these maps preserve complementation, which we do not want in cases of expansion and clarification. Instead, we can use a *lattice homomorphism* which preserves the meet and join structure but nothing else, allowing complementation relations and the bounds $\top$ and $\bot$ to vary between the two structures.

Now, if the mapping is to be an identity criterion for propositions, it needs to be one-to-one.[16] Putting this together, we are led to consider the *lattice embeddings*.

> **Lattice embedding.** A map $h: \mathcal{A} \to \mathcal{B}$, between two lattices $(\mathcal{A},\vee,\wedge)$ and $(\mathcal{B},\sqcap,\sqcup)$, is a lattice embedding iff it is a one-to-one lattice homomorphism. That is, a one-to-one map that is meet- and join-preserving: $\forall X, Y \in \mathcal{A}, h(X \vee Y) = h(X) \sqcup h(Y), h(X \wedge Y) = h(X) \sqcap h(Y)$.

A lattice embedding maps each proposition from the old algebra to a proposition in the new algebra, and preserves the lattice operations, meet and join.[17] It represents a subjective identification on the part of the agent; it is not intended to say anything about which propositions are really identical in some metaphysical sense. This will become important in resolving Mahtani's cases where I claim that the agent has a choice between two plausible identification schemes. In the definition above I used the notation just introduced to make $\mathbb{A}$ and $\mathbb{B}$'s operations explicit and distinct, but I won't be as careful with the notation from here on. Using an embedding means that there's no

---

[16] Is it question begging against Mahtani to assume that there is an identity relation between propositions in cases like hers? I do not think so: Mahtani's contention with the case is to show that Reverse Bayesianism cannot give a sensible answer, not to assert that there is no answer to which of BOB and OTHER is identical with TENANT. Indeed, Mahtani's intuitions for the right answer in this case support an identification of TENANT with BOB for the purpose of belief extension, as I discuss in §6.

[17] Regarding the order-theoretic definition of lattices: defined this way, lattice homomorphisms are order-preserving (Davey and Priestley 2002: Proposition 2.19, p.44).

need to differentiate between $\vee$ and $\sqcup$, or $\wedge$ and $\sqcap$: because an embedding is one-to-one, the image $h(\mathcal{A})$ is a sublattice of $\mathbb{B}$ which is isomorphic to $\mathbb{A}$.[18] Note that there is no guarantee that $h(\neg X) = \sim h(X)$, or that $h(\bigvee \mathcal{A})$ is the top of $\mathbb{B}$.

Let's apply this new machinery. First, consider Naledi's expansion experience. She begins with the partition {RAIN, SUN} in algebra $\mathcal{A}$, and ends with the partition {Rain, Sun, Snow} in $\mathcal{B}$. (Note the new capitalisation, as these are distinct objects.) The natural embedding $h: \mathcal{A} \to \mathcal{B}$ maps RAIN $\mapsto$ Rain and SUN $\mapsto$ Sun. The mapping of RAIN to Rain represents Naledi's identification of the two and reflects our judgement that the RAIN proposition is 'still there'. The fact that h is a lattice homomorphism guarantees that h(RAIN $\vee$ SUN) = Rain $\vee$ Sun and h(RAIN $\wedge$ SUN) = Rain $\wedge$ Sun. In this latter case, the conjunction of RAIN and SUN is a contradiction, RAIN $\wedge$ SUN = $\bigwedge \mathcal{A}$, and our mapping gets us this too as, in the new algebra, Rain $\wedge$ Sun = $\bigwedge \mathcal{B}$. However, in the initial algebra $\neg$RAIN = SUN, whereas in the new algebra $\sim$ Rain = Sun $\vee$ Snow. So $h(\neg$ RAIN) = h(SUN) = Sun $\neq \sim h$(RAIN). This tracks our intuitions about what has changed for Naledi.

In the case of Naledi's clarification, she begins with {RAIN, SUN} and ends with two distinct partitions {Rain, $\sim$Rain} and {Sun, $\sim$Sun}. Naledi identifies RAIN with Rain, and SUN with Sun. The embedding which does this maps RAIN$\wedge$SUN$\mapsto$Rain$\wedge$Sun, but this latter proposition is not the bottom of the new algebra, representing Naledi's new understanding that Rain and Sun are not contradictory propositions. Once again, we note that complementation is not preserved, since $h(\neg$RAIN$) = h$(SUN) = Sun $\neq \sim$ Rain $= \sim h$(RAIN).

We've now arrived at our recipe for modelling awareness growth.

> **Modelling recipe.** An agent's awareness state is modelled by a Boolean algebra $\mathbb{A} = (\mathcal{A}, \wedge, \vee, \neg)$. After their awareness grows, they have a new awareness state: the Boolean algebra $\mathbb{B} = (\mathcal{B}, \sqcap, \sqcup, \sim)$, where $\mathcal{B}$ contains the new propositions. We relate the old algebra to the new via a lattice embedding $h: \mathcal{A} \to \mathcal{B}$. The one-to-one association of propositions in $\mathcal{A}$

[18] A *sublattice* of a lattice $\mathbb{A}$ is a subset $\emptyset \neq \mathcal{C} \subseteq \mathcal{A}$, such that $\mathcal{C}$ is closed under the lattice operations: $\forall X, Y \in \mathcal{C}, X \vee Y \in \mathcal{C}, X \wedge Y \in \mathcal{C}$. It differs from a *sub-Boolean algebra*, which must contain the $\top$ and $\bot$ elements of $\mathbb{A}$.

with propositions in $\mathcal{B}$ ensures that the old propositions are 'in' the new algebra: for each $X \in \mathcal{A}$, $h(X) = x \in \mathcal{B}$.

The recipe does not tell us, at this stage, which embedding we need. Nor have I yet provided any substantive norms of awareness revision—we have simply observed the nature of the change and tried to represent it in an appropriate and precise framework.

5. Answering Mahtani's challenge

I will now apply this recipe to Mahtani's case of *The Other Tenant*. Recall that we start off with an initial awareness state consisting of LANDLORD, TENANT, and BOB. In a lattice theoretic presentation, we simply cannot represent BOB and TENANT as two different propositions if we are to capture the other things Mahtani says about them: that they are equivalent and that they therefore carry the same probability. At best we can say that these are two labels for the same proposition. There are two ways to see that we're forced into this choice: first, propositions are elements of the set $\mathcal{A}$ and sets contain only one copy of each item. Second, $\mathbb{A}$ is a Boolean algebra and so each element must have a *unique* complement—if TENANT and BOB have the same complement (LANDLORD) then they must be identical.

I don't mean to hide behind formalism here. The whole theory of probability is built on Boolean algebras, and these just aren't the kind of structures that can represent the kind of distinct-but-equivalent propositions that Mahtani gestures at. One could enhance the structure in some way, as others have done when considering the (hyper)intensionality of belief (for example, Chalmers 2011). I don't think that is required here. As I will show, we can capture the sense in which the original TENANT possibility 'splits' when your awareness grows in my framework. I leave it to proponents of a more complex model to show what mine is missing.

Propositional identity across awareness change is here modelled by the lattice embedding. What is puzzling about Mahtani's cases is that they offer us two choices for how to embed the original algebra into the new algebra. In Figure 2(a), I've shown the initial algebra $\mathbb{A}$ with the two possibilities that you start off being aware of: L for LANDLORD and TB for TENANT/BOB, where that double label reflects what I said above. Figures 2(b) and 2(c) show the new algebra $\mathbb{B}$, in which there's a proposition TENANT (t) with propositions BOB (b) and OTHER (o) entailing it. I've used lower-case

labels for these propositions in $\mathcal{B}$, because I want to insist that they're mathematically different entities, which will come to be associated with the elements of $\mathcal{A}$ via an embedding.

Here is one possible embedding, which I will label h: $\mathcal{A} \to \mathcal{B}$. h maps L to l, and TB to b. Under this embedding, the old possibility TENANT/BOB is mapped to the new possibility BOB. We can work out the rest from the fact that h is a lattice homomorphism. For the contradiction: $h(\bigwedge\mathcal{A}) = h(L \wedge TB) = h(L) \wedge h(TB) = l \wedge b = \bigwedge\mathcal{B}$. For the tautology: $h(\bigvee\mathcal{A}) = h(L \vee TB) = h(L) \vee h(TB) = l \vee b$, which is not the top element of $\mathcal{B}$. That is, of course, what we want: you previously thought that the landlord and Bob were the only possibilities for the singer, but then you became aware of the possibility of a second tenant. This also changes the complementation structure of the algebra: now you realize that, if it is not the landlord singing, then it might be Bob or the other tenant. The image of $\mathcal{A}$ in $\mathcal{B}$, under embedding h, is shown in Figure 2(b), indicated by the bolded letters in the lower left. This embedding makes *The Other Tenant* an example of expansion.

Another embedding, which I will label g, maps TENANT/BOB to the new possibility TENANT. It is shown in Figure 2(c), again via bolded letters. On this embedding, *The Other Tenant* is a case of refinement: where you previously thought in terms of 'the tenant' you now recognize two finer distinctions within this proposition.

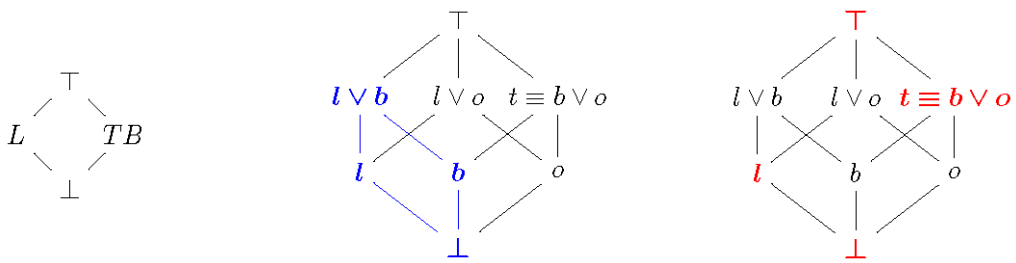

Figure 2. The Other Tenant. (a) The old algebra. (b) The new algebra, with the old embedded into the lower left bolded portion by h. (c) The new algebra, with a different, bolded, embedding g. On Mahtani's analysis, (b) is the preferred embedding for The Other Tenant.

5.1. A detour into belief extension

Mahtani introduced this case as a challenge to Reverse Bayesianism, and so we need to look at how my model feeds into that proposal in order to see whether we have made any progress. So let us reintroduce Reverse Bayesianism in the new formalism. The core definition is unchanged: when an agent undergoes awareness growth, their belief state should be rigidly extended to the new possibility space. However, we can immediately see that Bradley's definition of rigid extension, presented in §3, will not do: it assumed that a new algebra is formed by taking the (closure of the) union of the old set of propositions $\mathcal{X}$ with the new propositions $\mathcal{E}$. No questions about the identity of the 'old' propositions in the new algebra arose, because it was assumed to be obvious that they were the members of $\mathcal{X}$.

In redefining rigid extension, I will make the distinction between the new and old algebra explicit and relate them by a mapping. Let $\mathbb{A}$, $\mathbb{B}$, and h be defined as before. Let P be the probability function representing on $\mathcal{A}$ agent's prior credences and let P^+ be any probability function on $\mathcal{B}$. Then we have:

> **Rigid Extension (updated).** P^+ is a rigid extension of P to $\mathcal{B}$ iff for all $X \in \mathcal{A}, P^+\big(h(X)|h(\vee\mathcal{A})\big) = P(X)$.

From this we can derive a ratio rule for $X, Y \in \mathcal{A}$ with $P(X) > 0, P(Y) > 0$:

$$\frac{P(X)}{P(Y)} = \frac{P^+\big(h(X)|h(\vee\mathcal{A})\big)}{P^+\big(h(Y)|h(\vee\mathcal{A})\big)}$$

Back to *The Other Tenant*: Now that the definition of rigid extension explicitly tracks the two algebras and the mapping between them, we get quite different behaviour from Reverse Bayesianism in this case. We now see that there are two, quite sensible, Reverse Bayesian prescriptions corresponding to the two embeddings I introduced above. On h (Figure 2(b)) we get:

$$\frac{P(L)}{P(TB)} = \frac{P^+\big(h(L)|h(L \vee TB)\big)}{P^+\big(h(TB)|h(L \vee TB)\big)} = \frac{P^+(l|l \vee b)}{P^+(b|l \vee b)} = \frac{P^+(l)}{P^+(b)}$$

Whereas on g (red) we we get:

$$\frac{P(L)}{P(TB)} = \frac{P^+\big(g(L)|g(L \vee TB)\big)}{P^+\big(g(TB)|g(L \vee TB)\big)} = \frac{P^+(l)}{P^+(b \vee o)}$$

Notice that we don't ever get the situation Mahtani uses to criticize Bradley, where we are forced to assign TENANT and BOB the same probability, and therefore leave none for OTHER. That's because embeddings are one-to-one mappings, so there is no embedding which will map TENANT/BOB to TENANT and to BOB. In this model, the sense in which TENANT/BOB 'is' the proposition TENANT *and* 'is' the proposition BOB, is just that there exists an embedding on which it is mapped to TENANT and another embedding on which it is mapped to BOB. Our revised principle won't problematically assign the new possibility zero credence, in the way that Mahtani highlighted. This establishes my claim that the core of Mahtani's cases has nothing to do with Reverse Bayesianism as such, but rather depends on prior questions of awareness revision.

6. Choosing an embedding

We can now return to those prior questions and discuss how an agent should choose between potential identifications of propositions across awareness states—or, in the language of the model, how we select the right lattice embedding. Ideally, we want a general account, perhaps a set of adequacy conditions for an embedding that apply to all instances of awareness growth. I do not have such an answer at present. Instead, in this section I will outline a partial account, by examining a simple case with both a clear right answer and some features that help us identify what makes it the right answer. The case involves an agent whose relevant prior beliefs were set using the principle of indifference (POI). The POI, I claim, involves reasoning about the space of possibilities in a way that is prior to the assignment of credences. My special cases are situations in which this reasoning isn't disrupted by the awareness growth. My proposal is that what makes a choice of embedding correct is that it preserves reasoning the agent had already done about their initial space of possibilities and which they have no reason to revise. What makes these cases easy to work with is that this reasoning is easy to identify, in virtue of the credal assignment relying on it.

The good news is that this gives us insight into Mahtani's *Other Tenant*, since it is an example of such a case. The bad news is that the POI is a special case, and it isn't clear how to identify the ingredients of my proposal in other cases. We gain insight into the *kinds* of features awareness revision should depend on, but by example rather than general characterisation. We learn something else important, however. Reflecting on

cases where the agent must give up their prior possibilistic reasoning highlights a deep problem with Reverse Bayesianism, showing that it cannot be a fully general rule for belief extension.

Let's get to it. Consider an agent who assigns their prior credences by means of the principle of indifference. Mahtani discusses a second case that might be such an example:

> *The Other Tails.* An agent considers a UK 10 pence coin and wonders whether it will land HEADS or TAILS when tossed. They assign equal credence to these possibilities on the basis, we will suppose, of the principle of indifference. They then think about the image on the tails side of the coin. Initially they think all 10p coins have an image of a lion, so that TAILS and LION are equivalent propositions to them. Later, the agent becomes aware that some 10p coins have an image of Stonehenge. So, in their new awareness state, LION and STONEHENGE each entail TAILS. (Adapted from Mahtani 2021: 8984–85)

How should the agent's awareness and belief states change? Once again, we start with two labels, TAILS and LION, for what we will model as one proposition TAILS/LION. When the agent's awareness grows, they end up with a proposition TAILS covering two possibilities, LION and STONEHENGE. There are two ways to embed the old algebra into the new one. The first embedding, h, maps TAILS/LION to LION and is shown in the bolded letters in the lower left of Figure 3(b). The second, g, maps TAILS/LION to TAILS and is shown in Figure 3(c). Presenting this example in the context of Reverse Bayesianism, Mahtani says, quite reasonably, that the agent ought to keep the probabilities for HEADS and TAILS unchanged at 0.5. They have become aware of a new image on the tails side of the coin, which is intuitively irrelevant to these credences. Rendered in terms of the model, this intuition favours the g embedding, in a context where the agent goes on to extend their beliefs in the Reverse Bayesian way.

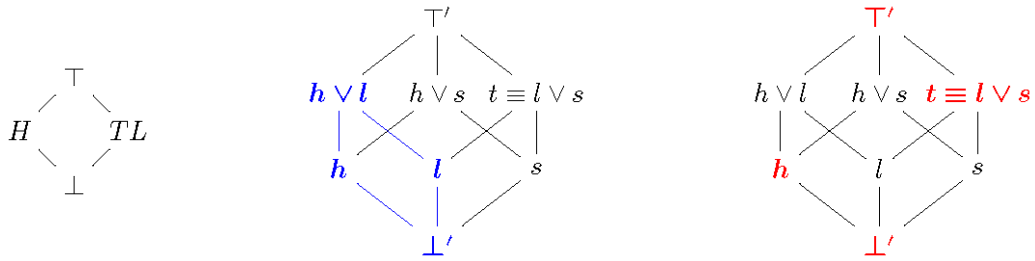

Figure 3. The Other Tails. (a) The old algebra. (b) The new algebra, with the old embedded into the bolded portion. (c) The new algebra, with a different, bolded, embedding. On Mahtani's analysis, (c) is the preferred embedding for The Other Tails.

Our goal is to understand why this is the right answer at the awareness revision stage, without relying on Reverse Bayesianism or our intuition that the right credal assignment is 0.5 for HEADS and TAILS. I stipulated that the agent relied on the POI, and it is here I will begin. This may not sound too promising, since it is a principle for the assignment of credence and, in modern treatments, is often justified on the basis of clearly belief-oriented considerations about evidence (White 2009) or epistemic utility (Pettigrew 2016). But it can help us get a handle on awareness revision when we reflect on how it operates.

To apply this principle, an agent identifies a partition of propositions as salient and judges that they have no reason to differentiate between the members of the partition. The partition members are then assigned equal credence. Notoriously, the choice of partition is critical: applying this procedure to different plausible partitions can result in contradictory assignments of credence. We need not worry here about whether this undermines the authority of the principle, but merely note that *if* an agent employed indifference reasoning, then they must have selected a partition as salient and made the requisite judgement about its members. This requires at least implicit reasoning about better and worse ways of carving up possibility space, and about the nature of and relations between the resulting possibilities. Indeed, as Hacking (1971) notes, indifference reasoning about subjective probabilities has historically been grounded in judgements of equipossibility, which often involves judgements that the possibilities are equivalent in some external sense—perhaps in light of a physical symmetry. This is crucial: the agent uses features of their context, such as their decision situation and properties of the objects they are reasoning about, to judge that certain propositions are

relevantly similar. Importantly, this reasoning takes place prior to the assignment of credences, and it concerns relations between propositions.

My proposal is that awareness revision should conserve this reasoning where possible. The motivation for this is as for conservative belief revision and indeed for Reverse Bayesianism, as discussed in §3: unnecessary changes are unjustified and risk error. Absent a reason to revise it, the agent should draw on their earlier reasoning and, in particular, on their judgement about which possibilities were salient and relevantly alike in their old awareness context. One might worry that these prior judgements are themselves rendered ambiguous by the awareness change, a worry I will dispel below. When faced with different ways they might identify old possibilities with new ones, the agent thus selects an identification which conforms to their prior reasoning about the old possibilities. Or, put in terms of my model, they choose an embedding which carries over judgements about the structure of the old algebra. I return to the subject of conservative awareness revision below, in §6.2.

In *The Other Tails*, when the agent initially applied the POI, they had two plausible partitions: one involves the two sides of the coin {HEADS, TAILS} and the other involves the two images {HEADS, LION}. In more typical cases of the POI these partitions would not map onto one another one-to-one. Nonetheless, we can ask ourselves which of these two partitions the agent had in mind when they applied the POI, by thinking about their reasoning about the space of possibilities. The context is that of a coin flip, which we know typically involves bets on which side lands up. The conventional names 'heads' and 'tails' are used to identify these sides, but often they have only a loose connection to the images actually found on the coin—for example, 20 and 50 cent Euro coins often do not have an image of a head on either side. The two sides are judged equivalent in virtue of presumed physical symmetries, which don't have anything to do with the images themselves. So, it is natural to suppose that the agent used the {HEADS, TAILS} partition, rather than {HEADS, LION}.

When the agent's awareness grows, they recognize that the new possibility (STONEHENGE) is an image-possibility, rather than a side-possibility. This involves recognising a similarity between LION and STONEHENGE, and the apprehension of their logical relation to TAILS. Thus, they grasp that their space of possibilities is as shown in Figure 3(b) and (c). They must now choose an embedding: decide to identify the old

TAILS/LION proposition (TL) with Lion (l) or with Tails (t). The agent now sees that their old TAILS concept was ambiguous— they previously took it to be equivalent to LION and, now that the two concepts have come apart, some of their prior reasoning in terms of TAILS/LION is ambiguous. But this ambiguity doesn't affect the sense of TAILS which mattered for their application of the POI, which was done in terms of the symmetry between the sides of the coin. The awareness growth experience doesn't alter that judgement of salience of equipossibility. They thus use this reasoning to guide how they identify propositions: TAILS/LION with Tails. Put in slightly different language, while one event had two labels, TAILS and LION, it was the TAILS label that was salient. Nothing about the awareness change experience prompts a revision of this reasoning. So, in order to preserve this reasoning, the agent selects the g embedding, shown in Figure 3(c).

If one were to now apply Reverse Bayesianism, this choice would mean that the agent retains 50:50 credence in HEADS and TAILS, as desired by both Mahtani and me. Of course, we could also get this result via a very different belief extension strategy: scrapping all prior credences and assigning new probabilities on the new algebra using the POI, applied to the partition the agent now finds most salient. I will return to this thought below, when I revisit the idea that an awareness change experience might prompt a revision of the agent's prior possibilistic reasoning, making it unsuitable as a basis for awareness revision.

6.1. The Other Tenant

We can now return to the original case, *The Other Tenant*. Once again, we start by checking the intuitively right answer and then reasoning to it independently. Recall our two options: the first embedding, h, mapped the TENANT/BOB proposition to BOB, and so feeding it into Reverse Bayesianism generates the requirement that there is no change in the relative probabilities assigned to the singer being the landlord or Bob. Whatever credence is assigned to the possibility that there is another tenant who is the singer, it needs to get its probability mass equally from that previously assigned to LANDLORD and TENANT/BOB. Recalling that Mahtani has the priors set up with $P(TB) = P(L) = 0.5$, this means that the extended probabilities for LANDLORD and BOB must be identical: $P^+(l) = P^+(b) = 0.5 - k$ and $P^+(o) = 2k$.

The g embedding, on the other hand, maps TENANT/BOB to TENANT. Reverse Bayesianism thus requires that the extended probabilities for LANDLORD and TENANT be identical. TENANT is entailed by BOB and OTHER, which are refinements of it. So, where you previously assigned equal credence to the singer being the landlord or the tenant (who you took to be Bob), now that you are aware of the possibility of another tenant, you assign equal credence to the singer being the landlord or *either* tenant: $P^+(l) = P^+(t) = 0.5$. There are no constraints on the credence assigned to BOB and OTHER so long as $P^+(b) + P^+(o) = 0.5$.

We now assess the case intuitively. I judge that the former assignment is more sensible, and Mahtani (2021, 8988) agrees: 'given that there might be two tenants, it is natural to suppose that your credence in TENANT should increase relative to LANDLORD'. Taking some credence solely from Bob to give to the other tenant seems unmotivated. (Note that this is the reverse of the situation in *The Other Tails*).

Mahtani doesn't tell us why you have equal credence in LANDLORD and TENANT/BOB. But we can test out the proposal above by supposing that, rather than having equally balanced evidence, you simply used the POI. To do so you would have had to choose a salient partition. There were two candidate partitions for the possible singers: {LANDLORD, BOB} and {LANDLORD, TENANT}. What distinguishes them is the use of the personal name Bob, which indicates that the former partition conceives of the possible singers as two people. The use of 'LANDLORD' as a label may be misleading here, so we might imagine you know the landlord's name is Ali, and momentarily relabel the two partitions {ALI, BOB} and {LANDLORD, TENANT}. Now the difference is clear, as is the reason for preferring the first partition. It is natural to think of possible singers in terms of their personhood, and unnatural to think of them in terms of their roles in a lease agreement. Let us check this by considering how the indifference reasoning might go: on the {ALI, BOB} way of thinking, the equipossibility reasoning is presumably that people are the relevant units of analysis, that there are two people in the house, and that the context (visiting a friend, showering in the morning) and available evidence (the singing) are naturally framed in terms of people. By contrast, the {LANDLORD, TENANT} partition would be salient in a situation involving these two roles in the legal agreement governing Bob's occupation of this apartment: rent negotiations, disputes about household obligations, or what have you. The evidence and the symmetries of this

situation fit more naturally with the person-based partition than the role-based partition.

Suppose that this is why you originally assigned equal credence to LANDLORD and BOB. When your awareness grows, you become aware of a new person-possibility. This does not give you reason to discard your initial judgement that the person-partition was salient and that its members were alike in a relevant sense. Thus, the embedding which identifies propositions on the basis of personhood is preferred to the embedding which identifies propositions on the basis of legal role. This is despite the fact that the TENANT/BOB proposition has been rendered ambiguous—you can reason about which element of its original content was relevant to you before, note that it is this part of the content which was used in the application of the POI, and then choose the propositional identification which allows you to preserve that reasoning. This supports the identification of TENANT/BOB with BOB. By choosing the corresponding embedding (h) and feeding it into Reverse Bayesianism we recover the intuitively correct solution.

6.2. Conservativity and Reverse Bayesianism

I have now shown how, in special cases, an agent should choose to identify propositions across changes of awareness, and shown that this is compatible with Reverse Bayesianism. Does that imply that an agent should always preserve their prior reasoning about how to partition the space of possibilities? No. I propose only that agents should do so when it is appropriate, and I will now explain how to assess that appropriateness. The logic underlying my proposal above is that of conservative change. This will be familiar to readers from the literature on belief revision, in both the qualitative (AGM) and quantitative (Bayesian conditioning) case. In that literature the key questions are 'conserve what?' and 'in what circumstances?'. A quick review of how these questions are answered in the case of Bayesian conditioning will help us answer the corresponding questions for my proposal.

Consider an ordinary Bayesian experience of learning, in which an agent gains some empirical information. Dietrich, List, and Bradley (2016) characterize the rationality of belief revision in response to such an experience with two conditions: Responsiveness and Conservatism. Responsiveness ensures that the final belief state matches the input from the experience. Conservatism ensures that the belief revision changes nothing that is not required to meet Responsiveness. In Bayesian learning, the agent learns a

proposition E. The axiom of Responsiveness thus demands that the posterior Q has $Q(E) = 1$. Conservatism demands that the belief revision preserves what the input is 'silent on'. The experience provides us with no reason to change what it is silent on. In the absence of such a reason, the thinking goes, we should make no change and instead retain our old belief. In the Bayesian case, this is codified by the Rigidity rule for conditional probabilities: $Q(\cdot\,|X) = P(\cdot\,|X), \forall X \in \Omega$. Dietrich, List, and Bradley (2016) show that Bayesian conditioning is the only way to satisfy Responsiveness and Conservatism for this kind of learning. They also characterize Jeffrey conditioning in this way and claim that all rational belief updates can be so represented. Each update rule (Bayes, Jeffrey, etc.) is characterised by a kind of input and what it is silent on.

I propose to think of awareness changes similarly. We have just seen that awareness growth experiences can also be silent on certain things. The *Other Coin* case gave the agent no reason to change their judgement of salience and equipossibility. My proposal above relied on this, as it motivated the agent's choice of an embedding. As we saw, awareness revision conducted in this way fitted neatly into belief extension by Reverse Bayesianism. But this is simply one kind of awareness growth experience. Just as there are non-Bayesian learning events, which don't obey the Rigidity condition, we should expect there to be awareness growth experiences which don't preserve these prior salience judgements.

Here is such a case. Recall Naledi, the South African wondering about the weather tomorrow in Stockholm. Let us suppose that she initially used the partition {WET, DRY}, taking these to be the relevant propositions for her choice of clothes. Suppose also that the assigned her credences using the POI. Upon becoming aware of SNOW, she faces a conundrum. Snow, she learns, is often wet, but sometimes not. In any case, it is not wet in the way that rain is. She could refine {WET, DRY} by considering the possibility of snow, but this seems to her to miss the point. Snow is a form of precipitation, and SNOW warrants a distinct clothing response from RAIN, a proposition which wasn't used in her reasoning previously but which she had the concept for. It is therefore more useful to switch to a precipitation partition and to regard SNOW as expanding that partition. In this case, the prior judgement about salience and equipossibility is disrupted. Indeed, the awareness revision process that I just sketched for Naledi is quite different from that described above. Here, it seems natural that Naledi must do some revisionary

reasoning about her old space of possibilities—switching from {WET, DRY} to {RAIN, SUN}—before considering how to relate old propositions to new ones.

The analogy I draw with Dietrich, List and Bradley's generalized belief revision is this: just as learning experiences come in kinds, which are characteristically silent on a part of the belief state and so facilitate conservative belief revision; so too are there kinds of awareness growth which are silent on aspects of the possibility space, and those facilitate conservative awareness revision. The kind of awareness growth I have discussed here are aspect on a particular aspect of the possibility space: salient partitions and prior judgements of equipossibility. The right form of awareness revision for this kind conserved these aspects of the agent's awareness state. And we saw that this, in turn, supported conservative belief extension, such as Reverse Bayesianism.

The WET/DRY example is not silent on this aspect of the awareness state: it requires a shift of the salient partition in the old algebra, a change that is so significant that it is no longer reasonable to hold on to the prior judgement of equipossibility. So, the right awareness revision doesn't conserve those aspects: it isn't mapping WET to some corresponding proposition in the new algebra. This is because it is responsive to the nature of the experience: in such a situation, Naledi should not try to preserve her prior reasoning about the space of possibilities.

Reverse Bayesianism is a proposal for conservative belief extension which assumes that the underlying awareness revision preserved the salient partition. The WET/DRY case is thus not one where Naledi ought to extend her beliefs using Reverse Bayesianism. She no longer judges those to be the salient propositions to consider, nor do they seem equipossible in her new environment: snow brings new ways for the world to be wet, and (fewer) new ways for the world to be dry. The Reverse Bayesian prescription seems wrong here, precisely because the awareness revision stage led to significant alterations in the space of awareness. When it comes to allocating credences to the propositions in her new algebra, she may as well begin again and do her indifference reasoning directly on {RAIN, SUN, SNOW}.

A similar problem can arise in cases of awareness growth by clarification, combined with credences assigned by indifference reasoning. This is because clarification splits a single partition {RAIN, SUN} into two independent partitions {Rain, ~Rain}, and {Sun, ~Sun}. The space of possibilities is so different that preserving prior judgements about

RAIN and SUN seems unwarranted. Previously the agent took them to be possibilities of the same kind, mutually exclusive members of a common salient partition. After the clarification, they take rain and sun to be independent features of the weather which can co-occur or not. What can we say about this case? There is an obvious intuition about awareness revision here: RAIN should map to Rain and SUN to Sun. But it can't be justified as I described above, in terms of preserving a prior judgement of salience of a (single) partition. I also judge that Reverse Bayesianism is the wrong approach to belief extension here: it makes no sense to hold on to a 1:1 ratio of credences between RAIN and SUN, if those were assigned on the basis of a (now rejected) judgement that they were members of a partition. This failure of Reverse Bayesianism is explained just as above: it assumes a certain form of awareness revision, which is not present here.

Note that once again this case is tied to the use of the POI. If the agent had some more substantial reasons for their credal assignments to RAIN and SUN, these may plausibly survive the shift. If, for example, they had seen a weather forecast assigning a 60% probability to RAIN, and if they have no reason to suppose that the weather forecaster made their reasoning error about the relationship between RAIN and SUN, then they can quite reasonably continue to defer to the forecast and assign a 60% probability to Rain. That is all quite distant from the discussion which has occurred here, and I raise it merely to emphasize the ways in which my discussion has been limited to a special case.

This all suggests that Reverse Bayesianism is only a plausible rule for belief extension in certain kinds of awareness growth. It seems to work in the kinds of special case I introduced above: where prior possibilistic reasoning survives and where it was the basis for the assignment of credences via the POI. I hypothesize that Reverse Bayesianism works more generally when the agent's change of awareness does not undermine the reasons for the prior credences which are candidates for extension. But as with my proposal for awareness revision, I do not have a general characterisation of 'reasons for prior credences' beyond the special cases I have examined. In any case, it is clear that Reverse Bayesianism is not a general strategy for belief extension following awareness growth.

Importantly, both the good and bad cases for Reverse Bayesianism that I have discussed are atypical. Each of these cases involved prior reasoning about the space of possibilities that was transparent and relevant to the assignment of credences. By

contrast, I assume that most of our contingent credences are what they are because of different pieces of evidence accumulated over time which we don't have conscious access to. It is also plausible that many of our credences don't depend so directly on our choice of a salient partition. Reverse Bayesianism may or may not work in other cases, but what we now know is that whether it is a sensible strategy for belief extension depends importantly on questions about the nature of the underlying awareness revision.

One way to develop my proposal beyond these limited cases is to turn to theories which link all beliefs to a privileged way of picking out possibilities. I have in mind recent work on belief as a question-sensitive attitude (Yalcin 2018; Hoek 2022, forthcoming). In this theory, all beliefs are understood as answers to (perhaps implicit) questions, and these questions in turn pick out (or are) partitions. It is possible that one could use the question that an agent had in mind prior to their awareness growth experience to motivate their choice of embedding in problem cases like Mahtani's. More radical awareness change experiences might cause changes of question, as in my example with Naledi and the {WET, DRY} partition above. I think that this is a promising avenue for investigation. Question-sensitive theories of belief need some development before they can really play this role, however. At present, the theory has focussed on qualitative belief and, in Hoek's case, was largely targeted at failures of logical reasoning. Hoek has some suggestions for how to frame credences as question sensitive, but no theory of updating question-sensitive credences. So, significant theoretical work would be required before this could be applied to my model for awareness revision and integrated with proposals for belief extension and (quantitative) belief revision.

7. Conclusion

Mahtani's cases highlight an important ambiguity in how awareness change has been hitherto discussed. They teach us that talking about 'the same proposition' in different awareness contexts is potentially problematic. In doing so, they identified a gap in prior discussions of awareness growth and proposals for belief extension, such as Bradley's presentation of Reverse Bayesianism.

I have made three contributions here: the first is clarifying that *this* is the import of those cases. Rather than being counterexamples to Reverse Bayesianism, they do not

target belief extension at all. Rather they put the spotlight on the prior, and previously neglected, stage of awareness revision. The second contribution is in providing a formal model for awareness revision, in the form of my lattice theoretic model in which propositional identification is represented by lattice embedding. The third contribution is my proposal for how agent's might rationally choose between different possible embeddings. The proposal is this: when the agent previously reasoned about the structure of the state of possibilities, and the awareness change experience is silent on that prior reasoning, then the awareness revision should be conservative with respect to it. I showed how this fixes the identification of propositions across awareness contexts in simple cases involving the principle of indifference.

There is clear room to develop this work. My proposal is a partial one and was here spelled out for a highly special and atypical case. My notion of 'reasoning about possibilities' needs development, as does the notion of conservative awareness change. Ideally, these would be neatly characterized in the same way that Dietrich, List, and Bradley (2016) do for belief revision. My hope is that we can characterise kinds of awareness change, identify what they are silent on, and then identify unique awareness revision procedures which fit each kind. This in turn would lead to new ways of researching belief extension and its relationship with the various kinds of awareness revision. But this will have to be wait until future work.

Acknowledgements

Thanks to Richard Bradley and Orri Stefánsson for discussions of and feedback on earlier drafts of this paper. Thanks also to audiences at the Higher Seminar in Theoretical Philosophy, Stockholm University; PROGIC 2022; and the WSC. Special thanks to Chloé de Canson for many early discussions on the topic of awareness growth as we each began working on it back in 2018-19.

Funding information

Riksbankens Jubileumsfond programme *Climate ethics and future generations*, grant M17-0372:1

Longview Philanthropy project *Rational Foundations of Longtermism*

ORCID iD

https://orcid.org/0000-0002-7663-8057

References

Bradley, Richard. 2017. *Decision Theory with a Human Face*. Cambridge University Press.

———. 2022. "Attitudes, Deliberation and Decisions." *Synthese* 200 (1): 2. https://doi.org/10.1007/s11229-022-03468-6.

Chalmers, David J. 2011. "Frege's Puzzle and the Objects of Credence." *Mind* 120 (479): 587–635. https://doi.org/10.1093/mind/fzr046.

Davey, B. A., and H. A. Priestley. 2002. *Introduction to Lattices and Order*. 2nd ed. Cambridge: Cambridge University Press.

Dekel, Eddie, Barton L. Lipman, and Aldo Rustichini. 1998. "Standard State-Space Models Preclude Unawareness." *Econometrica* 66 (1): 159–73. https://doi.org/10.2307/2998545.

Dietrich, Franz, Christian List, and Richard Bradley. 2016. "Belief Revision Generalized: A Joint Characterization of Bayes' and Jeffrey's Rules." *Journal of Economic Theory* 162: 352–71. https://doi.org/10.1016/j.jet.2015.11.006.

Earman, John. 1992. *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.

Fagin, Ronald, and Joseph Y. Halpern. 1987. "Belief, Awareness, and Limited Reasoning." *Artificial Intelligence* 34 (1): 39–76. https://doi.org/10.1016/0004-3702(87)90003-8.

Gillies, Donald. 2001. "Bayesianism and the Fixity of the Theoretical Framework." In *Foundations of Bayesianism*, edited by David Corfield and Jon Williamson, 363–79. Applied Logic Series. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-017-1586-7_15.

Hacking, Ian. 1971. "Equipossibility Theories of Probability." *The British Journal for the Philosophy of Science* 22 (4): 339–55. https://www.jstor.org/stable/686316.

Heifetz, Aviad, Martin Meier, and Burkhard C. Schipper. 2006. “Interactive Unawareness.” *Journal of Economic Theory* 130 (1): 78–94. https://doi.org/10.1016/j.jet.2005.02.007.

———. 2013. “Unawareness, Beliefs, and Speculative Trade.” *Games and Economic Behavior* 77 (1): 100–121. https://doi.org/10.1016/j.geb.2012.09.003.

Hill, Brian. 2010. “Awareness Dynamics.” *Journal of Philosophical Logic* 39: 113–37.

Hoek, Daniel. \{forthcoming\}. “Minimal Rationality and the Web of Questions.” In *Unstructured Content: Foundations and Applications of Possible Worlds Content*, edited by Dirk Kindermann, Peter van Elswyk, and Andy Egan. Oxford: Oxford University Press.

———. 2022. “Questions in Action.” *The Journal of Philosophy* 119 (3): 113–43. https://doi.org/10.5840/jphil202211938.

Karni, Edi, and Marie-Louise Vierø. 2013. “‘Reverse Bayesianism’: A Choice-Based Theory of Growing Awareness.” *American Economic Review* 103 (7): 2790–810. https://doi.org/10.1257/aer.103.7.2790.

Lakatos, Imre. 1968. “Changes in the Problem of Inductive Logic.” In *Imre Lakatos, Philosophical Papers*, edited by John Worrall and Gregory Currie. Vol. 2. Cambridge: Cambridge University Press.

Mahtani, Anna. 2021. “Awareness Growth and Dispositional Attitudes.” *Synthese* 198: 8981–97. https://doi.org/10.1007/s11229-020-02611-5.

Pettigrew, Richard. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Romeijn, Jan-Willem. 2005. “Theory Change and Bayesian Statistical Inference.” *Philosophy of Science* 72 (December): 1174–86.

———. 2012. “Conditioning and Interpretation Shifts.” *Studia Logica* 100: 583–606. https://doi.org/10.1007/s11225-012-9415-x.

Roussos, Joe. 2021. “Expert Deference as a Belief Revision Schema.” *Synthese* 199: 3457–84. https://doi.org/10.1007/s11229-020-02942-3.

Schipper, Burkhard C. 2015. “Awareness.” In *Handbook of Epistemic Logic*, edited by H. van Ditmarsch, J. H. Halpern, W. van der Hoek, and B. Kooi, 77–146. College Publications.

Schipper, Burkhard C. 2014. “Unawareness Gentle Introduction to Both the Literature and the Special Issue.” *Mathematical Social Sciences*, Unawareness, 70 (July): 1–9. https://doi.org/10.1016/j.mathsocsci.2014.03.002.

Shimony, A. 1970. “Scientific Inference.” In *The Nature and Function of Scientific Theories*, edited by R Colodny, 79–172. University of Pittsburgh Press.

Steele, Katie, and H. Orri Stefánsson. 2021a. “Belief Revision for Growing Awareness.” *Mind* 130 (4): 1207–32. https://doi.org/10.1093/mind/fzaa056.

———. 2021b. *Beyond Uncertainty*. Cambridge: Cambridge University Press.

Vallinder, Aron. 2018. “Bayesian Variations: Essays on the Structure, Object, and Dynamics of Credence.” PhD thesis, London: London School of Economics; Political Science.

Wenmackers, S, and Jan-Willem Romeijn. 2016. “New Theory About Old Evidence.” *Synthese* 193 (4): 1225–50.

White, Roger. 2009. “Evidential Symmetry and Mushy Credence.” In *Oxford Studies in Epistemology*, 3:161–86. Oxford: Oxford University Press.

Williamson, Jon. 2003. “Bayesianism and Language Change.” *Journal of Logic, Language and Information* 12 (1): 53–97. https://doi.org/10.1023/A:1021128011190.

Yalcin, Seth. 2018. “Belief as Question-Sensitive.” *Philosophy and Phenomenological Research* 97 (1): 23–47. https://doi.org/10.1111/phpr.12330.