



Awareness growth and dispositional attitudes

Anna Mahtani¹

Received: 31 July 2019 / Accepted: 26 February 2020
© The Author(s) 2020

Abstract

Richard Bradley and others endorse Reverse Bayesianism as the way to model awareness growth. I raise a problem for Reverse Bayesianism—at least for the general version that Bradley endorses—and argue that there is no plausible way to restrict the principle that will give us the right results. To get the right results, we need to pay attention to the (dispositional) attitudes that agents have towards propositions of which they are unaware. This raises more general questions about how awareness growth should be modelled.

Keywords Unawareness · Awareness growth · Bradley · Decision theory with a human face · Dispositions · Reverse Bayesianism

Classical Bayesian decision theory leaves no room for awareness growth. But in his book *Decision Theory with a Human Face* (Bradley 2017), Richard Bradley aims to give a decision theory for *real* agents, and it seems that real agents can grow in awareness. For example, I might be alerted to a possible state of affairs that had not occurred to me before; or I might discover that there is a further option available to me that I had not previously considered. Bradley gives an account of what rationality requires of agents without full awareness, and how rational agents' credal states can change when they grow in awareness. In this paper I raise a problem for this account, and argue that we need to model awareness growth quite differently, taking into account an agent's dispositional attitudes towards propositions of which (s)he is unaware.

Many thanks for numerous very helpful comments from Richard Bradley, Chloé De Canson, Orri Stefánsson, Marie-Louise Vierø, an audience at a workshop on Severe Uncertainty and Climate Policy at the Institute for Futures Studies in Stockholm, an audience at a workshop on Decision Theory & the Future of Artificial Intelligence at the ANU, and the anonymous referees for this journal.

✉ Anna Mahtani
A.Mahtani@lse.ac.uk

¹ Department of Philosophy, Logic and Scientific Method - Lakatos Building, London School of Economics and Political Science, Houghton St, London WC2A 2AE, UK

1 Bradley's account

Let us begin with Bradley's illustration of two different kinds of awareness growth. You are considering tomorrow's weather, and you are aware that it might be either cloudy, rainy, snowy or sunny, and also that it might be either cold or hot. Thus you have eight prospects to consider (as seen in the Table 1), and you have some credence in each of these prospects:¹

Table 1 Initial awareness state (weather example)

Cloudy		Rainy		Snowy		Sunny	
Cold	Hot	Cold	Hot	Cold	Hot	Cold	Hot

To begin with the first kind of case of awareness growth, which in the literature is called "refinement", let us suppose that you become aware of a further relevant category: the weather might be humid, or not humid. You now have sixteen prospects to consider (where "H" stands for "humid") (Table 2):

Table 2 Refinement (weather example)

Cloudy				Rainy				Snowy				Sunny			
Cold	Hot	H	¬H	Cold	Hot	H	¬H	Cold	Hot	H	¬H	Cold	Hot	H	¬H

This is a case where your increase in awareness has refined the categories that you were considering. Bradley writes:

"the mere fact that one's attitudes are defined on a domain that has proved too coarse does not give one any reason to change one's attitude to the coarse-grained prospects themselves. It follows that ... one's new degrees of belief... should agree with one's old ones over their common domain." (Bradley 2017, p. 257)

Thus for example, if you originally had a credence of 0.2 that it will be cold and cloudy, then after awareness growth your credence that it will be cold and cloudy should still be 0.2, with this 0.2 divided somehow between the possibility that it will be cold, cloudy and humid, and the possibility that it will be cold, cloudy and not humid.

Let us turn now to the second kind of case of awareness growth, which in the literature is called "expansion". Suppose that (from your state of awareness as given by Table 1) you become aware of the possibility that it might be misty, which is (let us suppose) incompatible with its being cloudy, rainy, snowy or sunny. This gives you ten prospects to consider (Table 3):

¹ Bradley uses the term "prospect" to mean "a possible way that the world might be" (Bradley 2017, p. 42).

Table 3 Expansion (weather example)

Cloudy		Rainy		Snowy		Sunny		Misty	
Cold	Hot	Cold	Hot	Cold	Hot	Cold	Hot	Cold	Hot

Here the number of possibilities that you need to consider has increased, but this time the new possibilities added are incompatible with the original possibilities. Bradley writes:

“...the key to conservative attitude change in cases where we become aware of prospects that are inconsistent with those that we previously took into consideration is that we should extend our relational attitudes to the new set in such a way as to conserve all prior relational beliefs and preferences.” (Bradley 2017, p. 257)

For example, suppose that when you were aware only of the 8 possibilities in Table 1, you had a credence of 0.4 that it would be rainy and cold, and a credence of 0.1 that it would be snowy and cold (and other credences in the other 6 possibilities so that your credences across all 8 summed to 1). Then your credence that it would be rainy and cold was 4 times as great as your credence that it would be snowy and cold. Once you become aware that it might be misty, your credences in the prospects in Table 1 will need to change to make room for some credence in the new possibilities of which you have become aware. But, on Bradley’s view, the proportions between the credences assigned to the prospects in Table 1 need to stay the same. Thus for example whatever your new credences in the possibility that it will be rainy and cold, and the possibility that it will be snowy and cold, your credence in the first needs to be 4 times as great as your credence in the second.

Thus Bradley has imposed certain restrictions on how a rational agent’s credal state develops in a case of awareness growth. Where the awareness growth is of the *expansion* kind, the rational agent will keep his or her credences (in the prospects of which she was previously aware) in the same proportions before and after the awareness growth. Where the awareness growth is of the *refinement* kind, then a rational agent’s credences (in the prospects of which she was previously aware) will remain the same after the awareness growth as before. Of course, if an agent has the very same credences in a set of prospects before and after awareness growth, then it follows that his or her credences in that set of prospects are in the same proportions before and after the awareness growth. And if we assume (as Bradley does) that a rational agent’s credences in all disjoint possibilities of which (s)he is aware must sum to 1 both before and after the awareness growth, then we can simplify by giving just one restriction for both kinds of awareness growth: the rational agent will keep his or her credences (in the prospects of which (s)he was previously aware) in the same proportions before and after the awareness growth.

We can follow Katie Steele and Orri Stefánsson and give this a more formal statement as follows (Steele and Stefánsson 2019; Steele and Stefánsson [Forthcoming](#)). In their definition, Steele and Stefánsson introduce the concept of a basic proposition—where a basic proposition does not involve any logical connectives. I follow

them in this because I want to relate the objections that I raise to Bradley's account to their account too. We let \mathbf{X} stand for the set of basic propositions of which an agent is aware. \mathbf{X} is the agent's 'awareness context' at that time. We say that an agent's awareness grows when the propositions of which s(he) is aware shifts from \mathbf{X} to \mathbf{X}^+ , where \mathbf{X} is a subset of \mathbf{X}^+ , and where \mathbf{X}_j is the (non-empty) set of propositions that are in \mathbf{X}^+ but not in \mathbf{X} . Then we can define the restriction described above as follows (Steele and Stefánsson, *Belief Revision for Growing Awareness* 2019):

Reverse Bayesianism *For any $A, B \in \mathbf{X}$ and according to any rational agent:*

$$\frac{P(A)}{P(B)} = \frac{P^+(A)}{P^+(B)}$$

This principle is entailed by the restrictions that Bradley imposes on rational agents in cases of awareness growth. The term 'Reverse Bayesianism' was coined by Edi Karni and Marie-Louise Vierø who pioneered this discussion. Reverse Bayesianism as a constraint on the credences over a state space—i.e. for states that are mutually exclusive—is endorsed by authors including Karni and Vierø (2013) and Wenmackers and Romeijn (2016). Steele and Stefánsson consider the principle as applied to a more complex proposition space, and Richard Bradley endorses the principle in full generality as (unlike Steele and Stefánsson) he does not restrict its application to basic propositions. I focus here on the principle as restricted to the set \mathbf{X} of all basic propositions—and the objection that I raise will apply to a less restricted version of the principle too.

Having set out the relevant version of Reverse Bayesianism, I turn in the next section to a problem for the account.

2 A problem for reverse Bayesianism

To explain the problem, I begin with an example that should prompt us to refine Reverse Bayesianism. I then set out another example that causes a problem for this refinement. I argue that there will be no good way of refining Reverse Bayesianism to avoid this problem.

2.1 Example 1: the other tenant

Suppose that you are staying at Bob's flat which he shares with his landlord. You know that Bob is a tenant, and that there is only one landlord, and that this landlord also lives in the flat. In the morning you hear singing coming from the shower room, and you try to work out from the sounds who the singer could be. At this point you have two relevant propositions that you consider possible, which I represent in Table 4, with LANDLORD standing for the possibility that the landlord is the singer, and BOB standing for the possibility that Bob is the singer:

Table 4 Initial awareness state (singing example)

LANDLORD	BOB
----------	-----

Because you know that Bob is a tenant in the flat, you also have a credence in the proposition (TENANT) that the singer is a tenant. Your credence in TENANT is the same as your credence in BOB, for given your state of awareness these two propositions are equivalent. Let us suppose, just for simplicity, that your credence in LANDLORD is 0.5 and your credence in TENANT (and so of course in BOB) is 0.5. We can represent your credences in the following table.

Table 5 Initial awareness state with added detail (singing example)

LANDLORD	TENANT BOB
0.5	0.5

Now let’s suppose that the possibility suddenly occurs to you that there might be another tenant living in the same flat, and that perhaps that is the person singing in the shower. Let’s assume that no other possibilities occur to you—e.g. it does not occur to you that it might be a visitor singing in the shower, or just a recording, or anything like that. Now the possibilities can be represented by the Table 6:

Table 6 New awareness state (singing example)

LANDLORD	TENANT	
	BOB	OTHER
?	?	?

We need to consider how your credence should be distributed once you are aware of all the possibilities outlined in Table 6. It is easy to see that Reverse Bayesianism will lead us into problems here. Both TENANT and BOB are possibilities that you were aware of before the awareness growth,² and you assigned them both the same credence (0.5). According to reverse Bayesianism, credences assigned to possibilities that you were aware of at the earlier time should remain in the same proportions after awareness growth has occurred. This means that TENANT and BOB need to stay in the same proportion—which means that they must be assigned the same value. Given that OTHER entails TENANT, and that BOB and OTHER are disjoint, it follows that

² It might be objected at this point that TENANT and BOB do not pick out the same propositions before and after the awareness growth, and so Reverse Bayesianism does not require your credence in TENANT and BOB to stay in the same proportion. There is a general question lurking here as to how we should identify the same propositions across different awareness contexts (a question that (Steele and Stefánsson 2019) also take up). The advocate of Reverse Bayesianism should be wary of claiming that any change in awareness context produces a new set of propositions, as this would drain Reverse Bayesianism of any content. We can take a proposition to be a set of possible worlds, but then there are difficult issues over which worlds exactly belong in the sets of possible worlds corresponding to TENANT and BOB before and after awareness growth in the example above. (Thanks to Richard Bradley and an anonymous reviewer for discussion on this point).

OTHER must be assigned zero.³ Thus it seems that in this case Reverse Bayesianism effectively rules out awareness growth: you can become aware of the new possibility of OTHER, but you must assign it zero credence.

This case looks rather different from the cases of awareness growth that have been mentioned elsewhere in the literature. In this example, we start out with different names (e.g. TENANT and BOB) for propositions which are equivalent relative to the possibilities (outlined in Table 5) that the agent is aware of at the start, though not equivalent relative to the possibilities (outlined in Table 6) that the agent ends up being aware of. Because of this, in this case of awareness growth we seem to have both a refinement and an expansion: we have a refinement relative to the possibilities LANDLORD and TENANT, but an expansion relative to the possibilities LANDLORD and BOB. Does this make the case somehow illegitimate? It certainly makes it different from the cases that are generally discussed, but nevertheless cases like this would be extremely common: often we might only be aware of one way that something could obtain, before discovering that there are many ways that that could obtain. We should expect any account of awareness growth to be able to handle this sort of case, and the fact that Reverse Bayesianism cannot is a shortcoming of that principle.⁴

Let us turn then to ways that we might refine the principle to avoid the problem. A first natural suggestion is that we should qualify the principle by stating that it holds only when *A* and *B* are disjoint propositions. Given that BOB and TENANT are not disjoint propositions, there would then be no requirement to keep the credences assigned to each in the same proportions. However the problem would still arise. For LANDLORD and TENANT are disjoint propositions, and so must remain in the same proportions as before—and given that they are exhaustive propositions they must therefore keep their original credence assignments of 0.5 and 0.5. LANDLORD and BOB are also disjoint propositions, and so must remain in the same proportions as before, which means (given that the credence assigned to LANDLORD is fixed at 0.5) that BOB must be assigned 0.5. Thus once again both TENANT and BOB are assigned the same credence of 0.5 and so OTHER must be assigned 0.

An alternative refinement from Steele and Stefánsson (2019) looks more promising at first sight. These authors also consider some counterexamples (of a rather different sort) to Reverse Bayesianism, and in response they impose a certain restriction on Reverse Bayesianism. The restriction states that the principle holds only where the

³ A defender of reverse Bayesianism might claim that here we should model belief revision in two steps: at the first step, awareness grows, and the new possibility (OTHER) receives zero credence and so reverse Bayesianism is satisfied; at the second step, the credence in the new possibility is updated. This does not solve the problem however, for if the new possibility is to receive a non-zero credence, then the second step is clearly not conditionalization. Many thanks to Orri Stefánsson for suggesting this point.

⁴ Steele and Stefánsson require that “an episode of growth in awareness be *either* a single expansion *or* a single refinement” (Steele and Stefánsson 2019, p. 13). They are not intending thereby to restrict the scope of their account to a certain sub-set of cases in which awareness grows only by expansion or restriction but not both. They write “[w]here an agent experiences awareness growth that is a mixture of expansion and refinement, this must be modelled as two sequential episodes of growth in awareness” (Steele and Stefánsson 2019, p. 14). But in the case of *The Other Tenant* it is not possible to separate out two events—one a growth of awareness by expansion, and the other by refinement. This kind of case could be characterised as a case of “partial refinement”, and Steele and Stefánsson’s account is not designed to cover such cases.

agent’s awareness growth is “evidentially irrelevant for A versus B ” (Steele and Stefánsson 2019, p. 23). They then define this evidential irrelevance condition as follows⁵:

Definition (*Evidential irrelevance*) For any $A, B \in \mathbf{X}$, we say that an agent’s awareness growth, from awareness context \mathbf{X} to \mathbf{X}^+ , where \mathbf{X}_j is the set of all basic propositions $X_j \in \mathbf{X}^+$ such that $X_j \notin \mathbf{X}$, is evidentially irrelevant for A versus B whenever:

$$\text{either } P^+ \left(A \mid \bigvee_{X_i \in \mathbf{X}_j} X_i \right) = 0 = P^+(B \mid \bigvee_{X_i \in \mathbf{X}_j} X_i)$$

$$\text{or } \frac{P^+(A \mid \bigvee_{X_i \in \mathbf{X}_j} X_i)}{P^+(B \mid \bigvee_{X_i \in \mathbf{X}_j} X_i)} = \frac{P(A)}{P(B)}$$

In other words, reverse Bayesianism only holds for propositions A and B where A and B meet one of two conditions: either they are incompatible with the disjunction of new propositions that the agent has become aware of; or conditionalizing on this disjunction of new propositions does nothing to change the agent’s relative credence in A and B .

This restricted version of reverse Bayesianism is not designed to handle the example of *The Other Tenant*, but let us in any case see what result we get when we apply it to this sort of example. In our example, the disjunction of new propositions that you become aware of is simply the proposition OTHER. We can see that this growth in awareness is evidentially irrelevant (given the definition) for LANDLORD versus BOB, because LANDLORD and BOB are both incompatible with OTHER:

$$P^+(\text{LANDLORD} \mid \text{OTHER}) = 0 = P^+(\text{BOB} \mid \text{OTHER})$$

Thus reverse Bayesianism can be applied here: the proportion of credence assigned to LANDLORD relative to BOB must not change when you become aware of OTHER.

But the other pairs of propositions do not meet the criteria. To illustrate this, take the propositions BOB and TENANT. It is not the case that BOB and TENANT are both incompatible with OTHER:

$$P^+(\text{BOB} \mid \text{OTHER}) = 0 \neq P^+(\text{TENANT} \mid \text{OTHER})$$

Nor is the second criterion met, for:

$$\frac{P(\text{BOB})}{P(\text{TENANT})} = \frac{0.5}{0.5} = 1$$

⁵ I have adapted the principle slightly for simplicity. In the text from (Steele and Stefánsson 2019, p. 23) instead of the first ‘ \mathbf{X} ’ in my quote they have ‘ $\mathcal{F}\mathbf{x}$ ’. The authors define a *possibility* as a truth-function, returning ‘true’ or ‘false’ for each basic proposition. A given basic proposition X_i can thus be associated with a set of possibilities—that is the set of all possibilities which return ‘true’ for that proposition. The authors then generate a Boolean algebra $\mathcal{F}\mathbf{x}$ from the sets of possibilities that correspond to the basic propositions in \mathbf{X} . This means that \mathbf{X} and $\mathcal{F}\mathbf{x}$ are not quite the same thing, but the difference does not matter for my purposes.

and yet:

$$\frac{P^+(\text{BOB}|\text{OTHER})}{P^+(\text{TENANT}|\text{OTHER})} = \frac{0}{1} = 0$$

and so:

$$\frac{P^+(\text{BOB}|\text{OTHER})}{P^+(\text{TENANT}|\text{OTHER})} \neq \frac{P(\text{BOB})}{P(\text{TENANT})}$$

Thus the awareness growth is not evidentially irrelevant for BOB versus TENANT, and so reverse Bayesianism need not be applied to this pair. Similarly the awareness growth is not evidentially irrelevant for LANDLORD versus TENANT. Thus the relative credence assigned across the pairs {BOB, TENANT} and [LANDLORD, TENANT] can vary before and after the growth in awareness.

This means that any distribution of credence following the pattern below is rational (for any $0 \leq k \leq 0.5$) (Table 7):

Table 7 Conforming to restricted reverse Bayesianism (singing example)

LANDLORD	TENANT	
	BOB	OTHER
$0.5 - k$	$0.5 - k$	$2k$

Just as before the awareness growth, you are assigning the same credence to LANDLORD as you are to BOB, but the amounts assigned to LANDLORD and BOB may be smaller so that there is some credence spare to assign to OTHER. Before your awareness growth, your credence in LANDLORD was the same as in TENANT—but now (for any $k > 0$) you have a greater credence in TENANT than in LANDLORD. This does not violate restricted reverse Bayesianism, because as we have seen your awareness growth is not evidentially irrelevant for LANDLORD versus TENANT. This seems like a good result, and it is in line with the way that we would naturally expect your credal state to change on becoming aware of OTHER: given that there might be two tenants, it is natural to suppose that your credence in TENANT should increase relative to LANDLORD.

The problem is that this will not work in every case. Here is an alternative case which has a similar form, but here we have quite different intuitions.

2.2 Example 2: the other tails

You know that I am holding a fair ten pence UK coin which I am about to toss. You have a credence of 0.5 that it will land HEADS, and a credence of 0.5 that it will land TAILS. You think that the tails side always shows an engraving of a lion. So you also naturally have a credence of 0.5 that (LION) it will land with the lion engraving

face-up: relative to your state of awareness, TAILS and LION are equivalent. We can represent your credal state with Table 8:

Table 8 Initial awareness state with added detail (coin example)	HEADS	TAILS LION
	0.5	0.5

Now let's suppose that you somehow become aware that occasionally ten pence coins have something other than a lion engraving on the tails side. In particular, you become aware that there are some ten pence coins that have an engraving of Stonehenge on the tails side. Let's assume that no other possibilities occur to you. Then the propositions of which you are aware are given in Table 9:

Table 9 New awareness state (coin example)	HEADS	TAILS	
		LION	STONEHENGE
	?	?	?

What should your credences be in these propositions? Formally the situation is just as it was for the previous example of *The Other Tenant*. There I showed that if you simply apply reverse Bayesianism without restriction, then your credence in BOB and TENANT would each need to remain at 0.5, leaving no room for OTHER. We can run an analogous argument here. You gave TAILS and LION the same credence before your awareness growth, and so by reverse Bayesianism you must give TAILS and LION the same credence after your awareness growth. But because LION and STONEHENGE are incompatible, and because STONEHENGE entails TAILS, this means that STONEHENGE must have a credence of zero. Thus given reverse Bayesianism, you can become aware of the proposition STONEHENGE, but you must give it a credence of zero. This is obviously not the account of awareness growth that we are looking for.

Let us then move to restricted Bayesianism, which seemed to solve our problems in the example of *The Other Tenant*. Can it also solve our problems in this example of *The Other Tails*? We get a superficially similar result: the growth in awareness is evidentially irrelevant for HEADS versus LION, because both HEADS and LION are incompatible with STONEHENGE. Thus restricted reverse Bayesianism requires that you keep your credence in HEADS and LION in the same proportion as before, which means that they must each receive the same credence as each other. In contrast, the growth in awareness is not evidentially irrelevant for HEADS versus TAILS: it is neither the case that HEADS and TAILS are both incompatible with STONEHENGE (for TAILS is not), nor that the (post awareness growth) proportion of your credence in HEADS conditional on STONEHENGE (0) divided by your credence in TAILS

conditional on STONEHENGE (1) is the same as the (pre-awareness growth) proportion of your credence in HEADS (0.5) divided by your credence in TAILS (0.5). Given that the growth in awareness is not evidentially irrelevant for HEADS versus TAILS, you are not required to keep your credence in HEADS and TAILS in the same proportion as before. In short, after your growth in awareness any distribution of credences following the pattern below is rational (for any $0 \leq k \leq 0.5$) (Table 10):

Table 10 Conforming to restricted reverse Bayesianism (coin example)

	HEADS	TAILS	
		LION	STONEHENGE
	$0.5 - k$	$0.5 - k$	$2k$

Thus after you become aware that the engraving on the ‘tails’ side of a ten pence coin is not always of a lion, but might alternatively be of Stonehenge, you can make room for some non-zero credence in the possibility STONEHENGE by reducing both your credence in HEADS and your credence in LION by some amount (keeping your credence in HEADS and LION the same as each other). This will mean that after your growth in awareness, your credence in HEADS is less than your credence in TAILS. This does not violate restricted reverse Bayesianism, because the growth in awareness is not evidentially irrelevant to HEADS versus TAILS.

The problem is that this entirely violates our intuitions about how a rational agent would adjust his or her credences on becoming aware of the possibility STONEHENGE. When you become aware that there is more than one type of engraving that might be on the tails side of a coin, that does not increase your credence in TAILS relative to HEADS. Your credence in HEADS and TAILS should remain at 0.5. Your credence in LION (but not in HEADS) ought to decrease to make room for some non-zero credence in STONEHENGE. In this case it ought to be your credences in HEADS and TAILS that remain in the same proportions—not your credences in HEADS and LION. Restricted reverse Bayesianism gives us the wrong result here, and this is reason enough to reject restricted reverse Bayesianism in its current form.

We might try to amend restricted reverse Bayesianism so that it gives us the right result here—that is, so that it requires that we keep the proportion of credence assigned to HEADS and to TAILS the same, but allows us to vary the proportion of credence assigned to HEADS and to LION. But it is not obvious what general restriction to place on reverse Bayesianism that will give us this result. Furthermore, we will not want this restriction to mischaracterise how a rational agent would adjust his or her credences in response to awareness growth in the previous example of *The Other Tenant*: there we want your credence in LANDLORD and BOB to remain in the same proportions—not your credence in LANDLORD and TENANT. The two examples are formally analogous, but we want different results in each. To see this, let us characterise {LANDLORD, TENANT} and {HEADS, TAILS} as the coarse-grained pairs, and {LANDLORD, BOB}, and {HEADS, LION} as the fine-grained pairs (of which you are aware from the start). In the case of *The Other Tails*, we want your credences in the coarse-grained pair {HEADS, TAILS} to remain in the same proportions after the

awareness growth, and to allow the proportion of your credences in the fine-grained pair {HEADS, LION} to vary. In contrast in the case of *The Other Tenant*, we want your credences in the fine-grained pair {LANDLORD, BOB} to remain in the same proportions, and to allow the proportion of your credences in the coarse-grained pair {LANDLORD, TENANT} to vary. There seems to be no formal set of restrictions that could give us this result. Thus the prospects for Reverse Bayesianism (at least, the version of it that Bradley endorses) seem bleak.

What has gone wrong here? I claim that the problem is that we have over-simplified the agent's earlier credal state. Our representation of the agent's credal state at the start is entirely silent about any attitude that the agent has towards OTHER and STONEHENGE. Of course, this is in line with the whole idea behind the literature on unawareness: on this view, agents do not have credences in propositions of which they are unaware. In the next section I argue that agents do have attitudes towards at least some propositions of which they are unaware, for they may have *dispositional* attitudes, and it is these dispositional attitudes that determine what happens when the agents grow in awareness.

3 The very idea of awareness growth

What is it to be unaware of a proposition? Bradley describes some different “senses or grades” of unawareness, and then writes: “What these situations of unawareness have in common is that certain contingencies or prospects are not available to the agent's consciousness at the time at which she is deliberating on some question” (Bradley 2017, p. 253). Bradley then states: “An agent can form attitudes only towards propositions that she is aware of” (Bradley 2017, p. 254).

Putting these two claims together, Bradley's claims entail that an agent can only form attitudes towards propositions that are available to his or her consciousness. It is hard to say exactly what it is for a proposition to be available to an agent's consciousness, but we can get the gist by thinking about the examples that Bradley discusses. Bradley gives a case where an agent is “unaware of the fact that there is a bus that goes to the town one wants to visit because one has never been there before”; a case where an agent has heard about something but fails to recall it; and a case where an agent has deliberately excluded a possibility from consideration (Bradley 2017, p. 253). In each of these cases there is some sort of cognitive barrier that prevents the agent from consciously considering the proposition. Bradley claims that in these sorts of cases, which we can characterise as cases where the relevant proposition is not available to the agent's consciousness, the agent cannot form any attitude towards the relevant proposition. I disagree. In both folk psychology and in more formal philosophy of mind, it is generally assumed that there are attitudes that an agent can take towards a proposition, even if (s)he does not (and even if, for various reasons, (s)he cannot) consciously consider that proposition. For example, we can say that an agent has a non-occurrent belief in a proposition while (s)he is asleep; that an agent knows something (the answer to a quiz question, perhaps) but can't recall it—but will recognise it as soon as it is mentioned; that an agent knows something deep down but is in denial; and so on. These attitudes cause (on some accounts in the philosophy of

mind) or consist of (on other accounts) various dispositions. For example, if an agent has a non-occurrent belief in some proposition, then (given various other mental and environmental conditions) the agent will assent to the proposition if someone else expresses it; if the agent knows something but can't recall it, then (s)he might state it if given a clue, and/or might be able to recall it at a later point; if an agent knows something deep down but is in denial, then (s)he might eventually come to admit the relevant proposition without gaining any new evidence.

I argue that in many examples from the literature on unawareness, the agent in question does have some attitude towards the relevant proposition—even if it is not available to his or her consciousness.⁶ The attitude may not manifest itself in any conscious thought, nor in any *actual* behaviour, but it will cause or consist of certain *dispositions*. To see one reason to accept this, let us return to the example of The Other Tenant. At the start of this scenario, you are aware of two relevant possibilities: LANDLORD and BOB. The possibility of OTHER is not something that has occurred to you, and so (let us assume) it is not something that is available to your consciousness. In the original scenario, at some point you became aware of the alternative possibility OTHER. Let's simplify the example by supposing instead that at some point you become abruptly aware not just that OTHER is a possibility, but that OTHER actually obtains. Suppose for example that you are in the kitchen running the tap and listening to the singing coming from the shower, wondering whether it is the landlord or Bob who is the singer (as these are the only two relevant possibilities available to your consciousness). By running the tap, you have disrupted the shower's water supply, and you hear the singer shout: "Hey Landlord—sort the water out! Bob and I are your tenants and we have rights!". You suddenly come to realise that OTHER is possible and indeed that it obtains. How do you react? That depends. We can fill out one version of the scenario as follows: this is the first time you've visited Bob and his landlord, they live in what looks like communal accommodation, and you just never thought to ask whether anyone else lives there too. In this version of the scenario, though the possibility of OTHER had not occurred to you before, it is not a great surprise to discover that OTHER obtains. Now let's drop that version of the scenario, and fill things out differently: you are very close friends with Bob and his landlord, you have stayed with them many times in their cosy homely flat which you know very well, and though they generally share all their news with you, they have never mentioned anything about another person living in their flat. In this case, the possibility of OTHER had never occurred to you before, and you are astonished to discover that OTHER obtains. Your level of surprise is very different in these two scenarios, and you might behave differently too—asking about the other tenant with mild curiosity in one case, and reacting with bewilderment in the other. What can explain this difference in your levels of surprise? On Bradley's view, there is no difference in your relation towards OTHER at the start of these two versions, for in both versions you simply have no attitude whatsoever towards this proposition. How then can we explain your difference in reaction in the two versions when you suddenly discover OTHER? On Bradley's

⁶ I claim that this holds for all the cases of unawareness that Bradley discusses. There may be cases for which this does not hold—perhaps cases involving radically new concepts, or transformative experiences (Paul 2014). These are controversial cases, and I don't attempt to explain these. Similarly, I don't attempt to apply this account to cases of childhood learning.

account there seems to be no good way to explain this difference: we would just have to take your level of surprise as a brute fact with no psychological explanation. This is unappealing. I propose that instead we should say that you do have an attitude towards OTHER at the start of the scenario, and your attitude towards OTHER is different in the two versions of the scenario: it is this difference in attitude that explains your different reactions on discovering that OTHER obtains.

What attitude should we say that you have towards OTHER at the start of the scenario (in either of the versions)? Should we say that it is the attitude of having some credence in OTHER? That depends on how we understand credences. One of the limitations of the Bayesian framework is that we end up representing an agent's epistemic state—in all its complexity—simply by a credence function. An agent can have all sorts of epistemic attitudes towards a proposition, and if we wished to fully characterise an agent's attitude towards a proposition, we'd need to think not just about how strongly an agent believed a proposition, but about many other dimensions too, such as how present the proposition is to the agent's consciousness. After all, an agent can be convinced of a proposition whether or not (s)he is currently considering it (as you no doubt were convinced a moment ago, and still are, that giraffes are mammals); and similarly a proposition of which an agent is unconvinced might equally well be something that (s)he is currently considering or not. The Bayesian framework does not allow us to represent all of these complexities, for on the Bayesian framework, an agent's attitude towards a proposition is represented by a single number. Thus on the Bayesian framework we do not get a full description of an agent's epistemic state, but rather a model of some aspect of that state. But what aspect? Here we have choices to make. We might choose to say that we are modelling conscious conviction—so that an agent who believes some proposition but is not currently conscious of it does not have a positive credence in that proposition. Or we might choose to say that we are modelling dispositional betting behaviour—but dispositions under what circumstances? Under circumstances in which the relevant proposition has been brought to the agent's attention by the offer of a bet? In that case an agent can have a credence in a proposition which (s)he is not consciously considering because of how (s)he *would* bet on that proposition were the proposition to be brought to his or her attention by the offer of a bet. And perhaps an agent can have a credence in some proposition even if there is some barrier to her consciously considering it—provided that that barrier is overcome by the agent being offered a bet on the proposition. Alternatively we might say that we are modelling dispositional behaviour more generally, in which case perhaps an agent can have a credence in some proposition even if (s)he could never become conscious of it, simply because our best explanation of his or her actual or dispositional behaviour involves the claim that (s)he has such a credence.

On many of these views, an agent can have a credence in a proposition even if those propositions are—for some reason or other—not available to the agent's consciousness. For example, to take one of Bradley's examples, even if you have not been to a town before and so you are unaware of the possibility that a bus goes there, you might still have some credence in that proposition, for you have dispositional betting behaviour (you would accept or reject various bets on the proposition were you to be offered them) and you also have more general relevant dispositions to behave in certain ways (you would react with astonishment—or not—on seeing a bus go past on

its way to the town; if you were to pass a bus stop, you might or might not consider it worth stopping to see if there are any buses going to the town, and so on). One option then is to say that in at least some cases of apparent unawareness, you had a credence in the relevant proposition all along. This is to deny the phenomenon of unawareness—at least for many of the examples that are put forward to motivate it. An alternative option would be to say that in these cases you do not have a credence in the relevant proposition, but nevertheless do have an attitude towards that proposition. For simplicity, let us characterise these two views together as claiming that you take attitudes not only towards propositions of which you are consciously aware, but also towards propositions of which you are not consciously aware.

On these views we can at least partially explain your level of surprise on discovering that OTHER obtains. Whether you are bewildered or just mildly surprised depends on the attitude that you took towards the proposition OTHER before you became consciously aware of it. If we think that you had a credence in this proposition all along, then presumably your credence was higher in the version where you are visiting Bob and his landlord for the first time in their communal accommodation, than in the version where you have visited Bob and his landlord multiple times in their cosy shared flat. And presumably your conditional credences were also different: your credence that Bob and his landlord have been deceiving you, conditional on OTHER, was much higher in the second version than in the first. Thus we can just use Bayesian Conditionalization to explain the differences in your level of surprise and your behaviour on discovering OTHER—provided that we assume that you had a credence in OTHER (and related propositions) all along. If instead we suppose that you started out with some attitude other than credence towards OTHER, then the explanation is less straightforward, but we nevertheless know where to start: it is with your earlier attitudes towards the relevant propositions—including propositions of which you were not consciously aware. Thus by assuming that you can have attitudes in propositions of which you are not consciously aware, we can explain what would otherwise be mysterious.

We can also see more clearly why reverse Bayesianism is a mistake. According to reverse Bayesianism, when a rational agent becomes consciously aware of some proposition, then his or her relative credences in the propositions of which (s)he was previously consciously aware are fixed by his or her previous relative credences in those same propositions. We have seen that this is not the case. In the scenario of *The Other Tenant*, you would rationally keep your credences in the two fine-grained propositions (LANDLORD, BOB) in the same proportion and vary the proportion of your credences in the two coarse-grained propositions (LANDLORD, TENANT). In contrast, in the scenario of *The Other Tails*, you would rationally keep your credences in the two coarse-grained propositions (HEADS, TAILS) in the same proportions, and vary the proportion of your credences in the two fine-grained propositions (HEADS, LION). This holds even though your credences in all the propositions of which you were consciously aware were parallel at the start of the scenario. Something else is needed to fix what happens to the proportions of your credences after the growth in awareness—something other than just your earlier credences in the propositions of which you are consciously aware. The obvious answer is that it is your earlier atti-

tudes towards the propositions of which you were not consciously aware: propositions involving OTHER in one case, and STONEHENGE in the other.

In explaining or justifying an agent's credences, then, we can make reference to both conscious and unconscious attitudes towards propositions. How does that help us explain how your credences evolve in the cases of *The Other Tenant* and *The Other Tails*? Here I only attempt to gesture towards an answer to this question. Let us suppose that in these cases, you had credences all along in OTHER and STONEHENGE, and when your awareness growth occurred you became conscious of these propositions. What prompted this awareness growth? One possibility is that it was some new evidence that prompted the awareness growth. For example, perhaps in the case of *The Other Tenant*, there was something about the tone of the singing that made you suddenly aware of the possibility that the person singing in the shower might be neither the landlord nor Bob, but some other tenant: this evidence could both bring the possibility of OTHER to consciousness, and also raise your credence in this proposition. Exactly how much it raises your credence—and what happens to your credence in the other propositions—can all be determined by conditionalization in the usual way. Presumably your original credence in TENANT conditional on this evidence that you acquire is greater than your original unconditional credence in TENANT—and this is why your credence in TENANT increases when you acquire this evidence from the tone of the singing in the shower. Similarly in *The Other Coin* some piece of evidence might bring the proposition STONEHENGE into your consciousness—for example perhaps you catch sight of a different coin with an engraving of Stonehenge on the tails side. Here presumably your original credence in TAILS conditional on the evidence that you acquire is the same as your original unconditional credence in TAILS, and this is why your credence in TAILS does not increase when you acquire this evidence. Thus where you have credences in all relevant propositions all along, and where your growth in awareness is prompted by some new evidence, your changes in credence can be fully explained by conditionalization in the usual way.

But sometimes your growth in awareness comes without any apparent gain in evidence at all, as in cases where you just remember something, or when some possibility suddenly occurs to you.⁷ This growth in awareness can happen without any change in credence at all, as propositions can move in and out of consciousness without your becoming any more or less convinced of them. But sometimes when you become aware of a proposition, your credence in that proposition and/or others may shift. You might become more convinced of a proposition when you are consciously considering it, or the opposite might happen as in cases of therapy where a subconscious but deeply held belief is brought to consciousness and then dismissed. What rule of rationality can explain this sort of change in credence—where no new evidence prompted the change? This is I think a general challenge for the Bayesian. There are many everyday cases where we change our credences without acquiring fresh evidence, and do not seem to be irrational (at least not in the natural sense) for doing so. Many of these cases do not involve awareness growth at all: a philosopher might be actively and consciously considering 3 propositions for some time before realising how they are

⁷ Or—to give a mixed case—the growth in awareness might be prompted by a gain in evidence, but the gain in evidence might not account for the shifts in credence that occur when your awareness grows.

related, and then once the relation is seen her credence in one or more of the propositions may change, without her acquiring any new evidence, and without any change in her awareness of the relevant propositions. The Bayesian might rule that such an agent is irrational—perhaps clarifying that ‘rationality’ here means ideal rationality—and then it would be natural for the Bayesian to also judge as irrational any agent whose credences change when (s)he becomes aware of a proposition without gaining any new evidence. Alternatively, where a theorist gives an account that allows a rational agent to change his or her credence without gaining any new evidence but by some other process (e.g. just thinking about things), we should expect this account to cover cases where an agent changes his or her credence by bringing some propositions to consciousness.

I have only gestured at the ways that we might try to explain or justify how an agent’s credences can change when (s)he grows in awareness, but my key point is that these explanations can make reference to both conscious and unconscious attitudes that the agent holds—and so to an agent’s attitudes towards propositions of which (s)he is unaware. Explanations that limit reference to an agent’s attitudes to those towards propositions of which s(he) is aware are impoverished, and as we have seen they cannot explain the relevant data: specifically, they cannot explain why a rational agent’s credences diverge in the cases of *The Other Tenant* and *The Other Tails*.

4 Conclusion

I have argued that Bradley’s account of awareness growth faces a problem: it involves the principle of Reverse Bayesianism, and this principle faces counterexamples. The solution is to recognise that an agent can have an attitude—perhaps a certain credence, but certainly an attitude of some sort—towards a proposition even if it is not available to his or her consciousness. Furthermore, the agent’s credal state after an episode of awareness growth can depend on these attitudes that (s)he had to the propositions before she became aware of them. This raises more general questions about how widespread awareness growth is, and how it should be modelled.

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

References

- Bradley, R. (2017). *Decision theory with a human face*. Cambridge: CUP.
- Karni, E., & Vierø, M. L. (2013). “Reverse Bayesianism”: A choice-based theory of growing awareness. *American Economic Review*, 103(7), 2790–2810.
- Paul, L. A. (2014). *Transformative experience*. Oxford: OUP.

Steele, K., & Stefánsson, O. (2019). Belief revision for growing awareness. *manuscript*.

Steele, K., & Stefánsson, O. (Forthcoming). *Severe uncertainty: Reasoning with unknown possibilities*. Cambridge: CUP.

Wenmackers, S., & Romeijn, J. (2016). New theory about old evidence. *Synthese*, 193(4), 1225–1250.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.