

ARTICLE

How Should your Beliefs Change When your Awareness Grows?

Richard Pettigrew 

University of Bristol, UK

Email: Richard.Pettigrew@bris.ac.uk

(Received 6 April 2022; revised 3 July 2022; accepted 12 August 2022)

Abstract

Epistemologists who study credences have a well-developed account of how you should change them when you learn new evidence; that is, when your body of evidence grows. What's more, they boast a diverse range of epistemic and pragmatic arguments that support that account. But they do not have a satisfactory account of when and how you should change your credences when you become aware of possibilities and propositions you have not entertained before; that is, when your awareness grows. In this paper, I consider the arguments for the credal epistemologist's account of how to respond to evidence, and I ask whether they can help us generate an account of how to respond to awareness growth. The results are surprising: the arguments that all support the same norms for responding to evidence growth support a number of different norms when they are applied to awareness growth. Some of these norms seem too weak, others too strong. I ask what we should conclude from this, and argue that our credal response to awareness growth is considerably less rigorously constrained than our credal response to new evidence.

Keywords: Bayesian epistemology; credences; awareness growth; rationality; accuracy; betting arguments

Epistemologists who study partial beliefs or credences – often known as Bayesian epistemologists – boast a well-developed account of how you should change them when you learn new evidence; that is, when your body of evidence grows. What's more, over the past century, they have provided a diverse range of epistemic and pragmatic arguments that support that account. But they do not have a satisfactory account of when and how you should change your credences when you become aware of possibilities and propositions you have not entertained before; that is, when your awareness grows. In this paper, I consider the arguments for the credal epistemologist's account of how to respond to evidence, and I ask whether they can help us generate an account of how to respond to awareness growth. The results are surprising: the arguments that all support the same norms for responding to evidence growth support a number of different norms when they are applied to awareness growth. Some of these norms seem too weak, others too strong. I ask what we should conclude from this, and argue that our credal

© The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

response to awareness growth is considerably less constrained than our credal response to new evidence.

Suppose you are considering the political commitments of your friend Jan. Currently, you categorize a person's political views using the concepts *liberal*, *centrist*, and *conservative*. So you assign credences to three possibilities: *Jan is a liberal*, *Jan is a centrist*, and *Jan is a conservative*. Your current credences in these possibilities are known as your *priors*. Now you learn for certain that Jan isn't a conservative. What should be your new or *posterior* credences in the possibilities? The standard diachronic credal norm of Bayesian Conditionalization says that your posterior credence in a proposition should be your prior credence in that proposition conditional on the evidence you've learned; that is, if *E* is your evidence and *X* is any proposition to which you assigned a prior credence, Bayesian Conditionalization demands that your posterior credence in *X* is the proportion of the prior credence you assigned to *E* that you also assigned to *X*; that is, it is the ratio of your prior credence in *X* & *E* to your prior credence in *E*. So, if your prior credences were these:

<i>Jan is a liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
3/12	2/12	7/12

And if your evidence is that Jan isn't a conservative, then your posterior credences should be these:

<i>Jan is a liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
3/5	2/5	0

But suppose that, rather than learning that Jan isn't a conservative, you learn no new evidence at all, and instead become aware of a proposition that you have not considered before and to which you do not assign a prior credence. Suppose, for instance, you learn no new evidence, but instead acquire a new concept; in particular, you acquire the concept *leftist*. You then use this concept to construct a new proposition, *Jan is a leftist*. You did not previously assign a credence to this proposition. What posterior credence should you assign to it? And what posterior credences should you assign to the other propositions, namely, those to which you did assign priors? For instance, if your priors were these:

<i>Jan is a liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
3/12	2/12	7/12

What should your posteriors be?

<i>Jan is a leftist</i>	<i>Jan is a liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
?	?	?	?

Or suppose I begin by categorising a person's political affiliation not as *liberal*, *centrist*, or *conservative*, but as *left* or *right*. So I assign priors to *Jan is on the left* and *Jan is*

on the right. But now I learn that, as well as this economic axis, there is another dimension along which a person’s politics can be categorised, namely, the social axis. So I acquire the concepts *authoritarian* and *libertarian*. I thereby become aware of the propositions *Jan is a left-libertarian*, *Jan is a right-libertarian*, *Jan is a left-authoritarian*, and *Jan is a right-authoritarian*. Again, the question arises: What posterior credences should I assign to these new propositions? And what posterior credences should I assign to the other propositions, namely, those to which I did assign priors? For instance, if your priors were these:

<i>Jan is on the left</i>	<i>Jan is on</i>
4/5	1/5

What should your posteriors be?

<i>Jan is left-libertarian</i>	<i>Jan is left-authoritarian</i>	<i>Jan is right-libertarian</i>	<i>Jan is right-authoritarian</i>
?	?	?	?

In the relatively small literature that treats cases like this, they tend to be called cases of *awareness growth*.¹ Typically, cases of awareness growth divide into two groups: *refinement cases* and *expansion cases*. In a refinement case, I refine a possibility or possibilities I previously considered. This is what happens, for instance, when I split *Jan is on the left* into *Jan is left-libertarian* and *Jan is left-authoritarian*, and split *Jan is on the right* into *Jan is right-libertarian* and *Jan is right-authoritarianism*. In an expansion case, on the other hand, I learn of a new possibility that is not a refinement of any possibility I have considered before. This happens when add a fourth possibility, namely, *Jan is a leftist* to the three I previously considered, namely, *Jan is liberal*, *Jan is a centrist*, and *Jan is conservative*.

So our question is this: what norms, if any, govern the relationship between the credences I assign before the awareness growth and the credences I assign afterwards?

1. Terminology

Before we begin, it will be helpful to introduce a little terminology. We represent your credal state at a given time by your *credence function*. This is a mathematical function c that takes each proposition X to which you assign a credence at that time and returns the credence $c(X)$ that you assign. We call the set \mathcal{F} containing all of the propositions to which you assign a credence your *agenda*. And we represent your credences on a scale from 0, which is minimal credence or 0%, to 1, which is maximal credence or 100%. So, if your agenda at a given time is the set of propositions \mathcal{F} , if X is in \mathcal{F} , and if c is your credence function, then $c(X)$ is at least 0 and at most 1.

Throughout, I will assume that your credence function is probabilistic. Suppose \mathcal{F} is a Boolean algebra of propositions, so that it includes $\neg X$ whenever it contains X and it contains $X \vee Y$ and $X \& Y$ whenever it contains X and Y . Then c is probabilistic so long

¹For a state-of-the-art philosophical treatment of these cases, see Steele and Stefánsson (2021). For other treatments, see Karni and Vierø (2013), Bradley (2017), Roussos (2020), Mahtani (2021), and Canson (Ms).

as (i) it assigns credence 1 to all tautologies and credence 0 to all contradictions and (ii) the credence it assigns to a disjunction of two contradictory propositions is the sum of the credences it assigns to the disjuncts. Now suppose \mathcal{F} is not a Boolean algebra. Then c is probabilistic so long as it is possible to extend c to a probabilistic credence function on a Boolean algebra that contains \mathcal{F} as a subset.

Bayesian Conditionalization says that, if c defined on \mathcal{F} is your credence function before you learn evidence E , and c' defined on \mathcal{F} is your credence function after you learn it, and if $c(E) > 0$, then it ought to be that $c'(X) = c(X|E) = \frac{c(X \& E)}{c(E)}$ for all X in \mathcal{F} .

Knowing different readers favour different levels of formalization, in what follows, I will not give any mathematical or symbolic presentation of a point without first giving it informally.

2. Impermissivism and awareness growth

Let me clear the ground for our discussion by noting that there are certain Bayesians who have ready-made answers to the problem of awareness growth. These are the proponents of impermissivism or the Uniqueness Thesis (Kocec and Titelbaum 2016). For them, there are no fundamental diachronic norms; not in the case of awareness growth and not in the case of evidential growth. Instead, at the fundamental level, there are just synchronic norms, but they are strong enough to determine exactly which credences you should have at any point in your epistemic life, and therefore at any point just after your awareness has grown. For instance, one brand of impermissivist says that your credences should match the so-called evidential probabilities conditional on your total evidence (Williamson 2000); another says your credences should maximize Shannon entropy among those that respect your total evidence, where Shannon entropy is a mathematical function taken to measure how unopinionated your credences are (Paris and Vencovská 1990; Jaynes 2003; Williamson 2010). Either way, the posterior credences you should have are wholly determined without reference to your prior credences.

3. Reverse Bayesianism and its discontents

Let's suppose, then, that we are not impermissivists about credal rationality – we are, instead, permissivists. What diachronic norms might we then impose on our credal response to awareness growth?

The most promising such norm is Reverse Bayesian. The standard diachronic credal norm, Bayesian Conditionalization, says that, upon learning a proposition with certainty, your posterior credences should be your prior credences conditional on that proposition. This is equivalent to demanding that, if you learn a proposition with certainty, your new posterior in that proposition should be 1, and the ratio between your posterior credences in two propositions that each entail that evidence should be the same as the ratio between your prior credences in those two propositions. Reverse Bayesianism says that, if your awareness grows so that you have to assign posterior credences to new propositions of which you've just become aware as well as to the old ones to which you assigned priors, and if you learn no new evidence, then, for certain pairs of propositions, the ratio between your posterior credences in them should be the same as the ratio between your prior credences in them. Which pairs of propositions? Any two

drawn from a certain subset of \mathcal{F} , the set of propositions to which you assigned credences before your awareness grew. We'll call this subset \mathcal{F}^\dagger . Keeping \mathcal{F}^\dagger unspecified for the moment, here is the schematic version of Reverse Bayesianism in symbols:

Reverse Bayesianism (RB) Suppose

- c defined on \mathcal{F} is your credence function at t ,
- c' defined on \mathcal{F}' is your credence function at t' ,
- $\mathcal{F} \subseteq \mathcal{F}'$, and
- between t and t' , the only epistemically relevant thing that happens to you is that you become aware of the propositions in \mathcal{F}' that aren't in \mathcal{F} .

Then, for all X, Y in $\mathcal{F}^\dagger \subseteq \mathcal{F}$, it should be that

$$\frac{c(X)}{c(Y)} = \frac{c'(X)}{c'(Y)}$$

We obtain different versions by specifying different subsets \mathcal{F}^\dagger of the prior agenda \mathcal{F} . We'll consider two: first, the version that Steele and Stefánsson (2021) consider in their book; second, a version closer to the original formulation by Karni and Vierø (2013).²

For Steele and Stefánsson, \mathcal{F}^\dagger is the set of propositions from \mathcal{F} that are *basic* in \mathcal{F} , where a proposition is basic in \mathcal{F} if it does not contain any other proposition in \mathcal{F} as a logical component. For instance, if $\mathcal{F} = \{X, Y, \neg X, X \rightarrow Y\}$, then X and Y are basic in \mathcal{F} , but $\neg X$ and $X \rightarrow Y$ are not.

Anna Mahtani (2021) provides a counterexample to this version of Reverse Bayesianism. You are staying with a friend, Bob, and while alone in the kitchen you hear someone singing in the shower. You assign credences to four propositions:

- *Landlord*, which says it's the landlord of the flat who is singing,
- *Tenant*, which says it's a tenant of the flat,
- *Bob*, which says that Bob is the singer, and
- $Bob \rightarrow Tenant$, which says if Bob is the singer, then the singer is a tenant.

You assign credence $\frac{1}{2}$ to *Landlord* and $\frac{1}{2}$ to *Bob*. Knowing Bob is a tenant, you assign credence 1 to $Bob \rightarrow Tenant$. And knowing no-one can be both landlord and tenant, you assign $\frac{1}{2}$ to *Tenant*. But now it occurs to you that there might be another tenant. You thus become aware of two further propositions to which you will now need to assign credences:

- *Other*, which says that it is someone other than Bob or the Landlord singing; and
- $Other \rightarrow Tenant$, which says that the other singer is a tenant.

You're certain that there is only one landlord, so you assign credence 1 to $Other \rightarrow Tenant$. What do you assign to *Other*? According to Reverse Bayesianism, the ratios between your posterior credences in *Landlord*, *Tenant*, and *Bob* should be the same as

²Richard Bradley (2017) proposes a closely related principle, which he calls *Rigid Extension*. The counterexamples to the two versions of Reverse Bayesianism that I consider are equally counterexamples to that.

the ratios between your prior credences in them, since these are the basic propositions of \mathcal{F} . But that entails that your posterior credence in *Bob* must equal your new credence in *Tenant*, since that was the case for your priors. And thus, if your posterior credences are to be probabilistic and they assign credence 1 to *Other* \rightarrow *Tenant*, they must assign credence 0 to *Other*. And that doesn't seem right. Mahtani concludes that this version of Reverse Bayesianism is wrong, and I agree.

Next, a version of Reverse Bayesianism that is closer to Karni and Vierø's original. Here, if \mathcal{F} contains some propositions that are pairwise disjoint and for each of which there is no stronger proposition in \mathcal{F} , then we let \mathcal{F}^\dagger be the largest set of such propositions. For instance, \mathcal{F} might include only the three possibilities, *Jan is a liberal*, *Jan is a centrist*, and *Jan is a conservative*, and \mathcal{F}' might then divide the first into *Jan is an ordoliberal* and *Jan is a classical liberal*, while leaving the second and third untouched. In this case, $\mathcal{F}^\dagger = \mathcal{F}$. Now suppose you initially assign equal credences of 1/3 to the three possibilities in \mathcal{F} . That is,

<i>Jan is a liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
1/3	1/3	1/3

Now suppose you learn that liberals in fact divide further into ordoliberals and classical liberals. Having discovered the logical space of political views includes four positions rather than three, you might quite reasonably wish to assign equal credences of 1/4 to each. That is,

<i>Jan is an ordoliberal</i>	<i>Jan is a classical liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
1/4	1/4	1/4	1/4

For instance, you might reason as follows: 'When I set my initial priors over the original three possibilities, I had no information about the prevalence of the three political positions among people in my society nor any information about their prevalence among people relevantly like Jan nor even any information about what properties of Jan would be relevant. So, I divided my credences equally between them. Now, having learned that there are in fact four such positions and still having no relevant information about their prevalence, I want to divide my credences equally between these four possibilities. I want to do this because it is how I would have assigned my priors if I'd been aware of the larger range of possibilities from the outset.'³

However, if you respond to the awareness growth in this way, you thereby violate the version of Reverse Bayesianism that we're currently considering. After all, *Jan is a liberal* and *Jan is a conservative* are both in \mathcal{F}^\dagger , and yet your prior credences in them are equal, while your posteriors are not, since your posterior credence that Jan is a liberal is the

³You might respond that *ordoliberal* and *classical liberal* are not categories at the same level of classification as *liberal*, *centrist*, and *conservative*. The latter are categories, perhaps, while the former are subcategories of the first of them in much the same way as *Panthera* and *Neofelis* are genera, while *Panthera leo* and *Panthera onca* are species within the former. But, at least in the case of political categorizations, it is not at all clear that we should respect the standard hierarchy of levels of categorization when we set our credences. For of course it can be a politically astute move to have your political position accepted as belonging to one of the higher levels in this categorization. So those levels should not be taken to indicate anything about the prevalence of those positions that belong to them.

sum of your posterior credence that she's an ordoliberal and your posterior credence she's a classical liberal, and each of those is equal to your posterior credence she's conservative.

Mahtani's example and mine share the same structure. In each, the awareness growth leads you to divide certain possibilities you had previously considered into more fine-grained possibilities, but it does not divide each of the original possibilities into the same number of new possibilities. And each shows that it is rational for you to violate a particular version of Reverse Bayesianism. For these reasons, I think Reverse Bayesianism must be wrong.

4. Looking to the arguments

When we seek to extend an existing norm – such as Bayesian Conditionalization – to cover a broader range of cases than it currently governs – such as the cases of awareness growth – there are a number of ways to proceed. We might simply consult our intuitions about the new cases and try to think of a general norm that captures all of those intuitions; we might think about the intuitive, informal motivation for the original norm and ask what that motivates in the new cases; or we might think about the more formal, philosophical arguments for the original norm and ask what happens when you apply them to these new cases. In this section, I'd like to do the latter and extend the existing arguments for Bayesian Conditionalization so that they cover not only cases in which you respond to new substantial evidence about the world, but also to cases in which your awareness grows.

Recall: Bayesian Conditionalization says that, if between an earlier and a later time, the only epistemically relevant thing that happens to you is that you learn a proposition with certainty, and if your prior credence function at the earlier time gave positive credence to that proposition, then your posterior credence at the later time should be obtained from your prior by conditioning on the proposition; that is, your posterior credences should be your prior credences conditional on the evidence you've acquired. In symbols:

Bayesian Conditionalization (BC) Suppose:

- (i) c defined on \mathcal{F} is your credence function at t ;
- (ii) c' defined on \mathcal{F} is your credence function at t' ;
- (iii) between t and t' , the only epistemically relevant thing that happens to you is that you learn proposition E with certainty.

Then, if $c(E) > 0$, then it should be that, for all X in \mathcal{F} ,

$$c'(X) = c(X|E) =_{\text{df.}} \frac{c(X \& E)}{c(E)}$$

I'll consider six arguments for this norm. They differ from each other along three dimensions. First: some argue directly that you should update by conditioning on your evidence, while others argue first that you should plan to update by conditioning on your evidence and second that you should follow through on any such updating plans. Second: some argue for a narrow scope norm, while others argue for a wide scope norm. Third: some argue for the norm by appealing to pragmatic considerations and some appeal to purely epistemic considerations.

Applied to the case of evidential growth, all of these arguments support Bayesian Conditionalization. But, when applied to the case of awareness growth, different arguments support different norms. Indeed, there are three norms supported by these arguments: *Anything Goes*, which is the trivial norm that places no constraints on your posterior; *Weak Reflection Principle*, which says that your prior places some reasonably modest constraints; and *Conservatism*, which says your prior places implausibly strict constraints.

4.1. Planning + narrow scope + pragmatic argument

The first argument for Bayesian Conditionalization is due to Peter M. Brown (1976). To present it, I must first set up some machinery. As above, let \mathcal{F} be the set of propositions to which you assign credences at the earlier and later time. As in all the arguments we'll consider, we assume that \mathcal{F} is finite. Now let $\mathcal{W}_{\mathcal{F}}$ be the set of classically consistent assignments of truth values to the propositions in \mathcal{F} . We might think of these as possible worlds grained only as finely as is needed to determine the truth value of each proposition in \mathcal{F} . We call them the possible worlds relative to \mathcal{F} . Brown makes the further assumption that, for each w in $\mathcal{W}_{\mathcal{F}}$, there is a proposition that is true at w and only at w . We might think of this as the state description of this world. We abuse notation and write this as w as well.

Brown's argument then assumes that there is some set $\mathcal{E} \subseteq \mathcal{F}$ of propositions from which the evidence you acquire between the earlier and later time will come; and he assumes that \mathcal{E} is a partition of logical space. That is, for each world w in $\mathcal{W}_{\mathcal{F}}$, there is exactly one E in \mathcal{E} that is true at w . An *updating plan* on \mathcal{E} is then a function that takes a proposition E in \mathcal{E} and returns the posterior credence function r_E the plan asks you to adopt should you learn E with certainty. We say that an updating plan r on \mathcal{E} is a *conditionalizing plan on \mathcal{E} for your prior c* if, for each element E of \mathcal{E} to which c assigns positive credence, the plan tells you to respond to learning it by conditioning c on it. That is, if $c(E) > 0$, then $r_E(-) = c(-|E) = \frac{c(- \& E)}{c(E)}$. Note: there can be more than one conditionalizing plan for a given prior, since a conditionalizing plan may recommend any posterior it likes as its response to a piece of evidence to which the prior assigns zero credence.

Next, imagine that, at the later time, after you've updated your credences in response to the evidence you've acquired, you will face a decision between a number of options. Let \mathcal{A} be the set of options between which you must choose. Given an option a in \mathcal{A} and a possible world w in $\mathcal{W}_{\mathcal{F}}$, we write $u(a, w)$ for the pragmatic utility you assign to the outcome of choosing a if w is the actual world. Brown assumes that you will choose the option with the greatest expected pragmatic utility, where that expected utility is calculated from the point of view of the credences you'll have at that time, which are of course the posteriors that your updating plan recommends. We can therefore assign pragmatic utilities not only to the options at a possible world, but also to an updating plan at a possible world. After all, given a world w , there is a unique E_w in \mathcal{E} that you will learn if w is the actual world; it is the unique proposition from \mathcal{E} that is true at w . And an updating plan r will then require you to adopt r_{E_w} as your posterior. And then you will be required to choose from \mathcal{A} the option $a^{r_{E_w}}$ that boasts maximum expected utility from the point of view of r_{E_w} . So we can take the pragmatic utility of r at w to be the utility of $a^{r_{E_w}}$ at w . That is, $u(r, w) = u(a^{r_{E_w}}, w)$. And we can then calculate the expected pragmatic utility of an updating plan r from the point of view of a prior c :

$$\text{Exp}_c(u(r)) = \sum_{w \in \mathcal{W}_F} c(w)u(r, w) = \sum_{w \in \mathcal{W}_F} c(w)u(a^{rE_w}, w)$$

Brown then shows that:

- (i) for any updating plan that is a conditionalizing plan for your prior and any decision problem you might face, the updating plan maximizes pragmatic expected utility from the point of view of your prior; and
- (ii) for any updating plan that isn't a conditionalizing plan for your prior, there is some decision problem you might face for which this plan does not maximize the expected utility from the point of view of your prior.

He takes this to establish that you should plan to update by conditioning. To establish Bayesian Conditionalization, we must add something more. We must add something like the norm that Sarah Paul calls *Diachronic Contenance*, which says that, if you intend to do something φ at a later time conditional on something else E happening in the interim, and if you acquire no reason to abandon that intention in the interim, then, if E does happen, you are rationally required to φ (Paul 2014).⁴ Then Bayesian Conditionalization follows.

How might we adapt this argument to cover the case in which it is not your evidence that grows but your awareness? There are two potential obstacles. As we will see, only one is surmountable.

First, what is an updating plan in this situation? When you anticipate learning new evidence between the two times, your updating plan is a function defined on the set of possible pieces of evidence you might acquire. When we think about awareness growth rather than evidential growth, your updating plan must be a function defined on the set of different ways your awareness might grow between the earlier and later times. So far, so good. But what does this formal account of the updating plan as a mathematical function represent in you as the individual with the plan? In the evidential case, we take it to represent something like a conditional plan or commitment or intention; that is, a plan or commitment or intention to do something – i.e., update in a particular way – conditional on something else happening – i.e., you receive a certain piece of evidence. And, at least implicitly, we take this plan to be something you might yourself consciously formulate and choose to adopt. But we can't do that in the case of awareness growth. In the case Brown considers, you can know the different possible pieces of evidence you might receive before you receive any of them. But you can't do that with awareness growth. As soon as I consider a particular way in which my awareness might grow, it immediately grows in exactly that way. I can't sit and wonder how I'd respond were I to consider the possibility of a second tenant in Bob's flat, since by doing so I already consider that possibility. I can't sit and wonder how I'd distribute my credences were I to learn the concept of a political leftist and formulate the proposition that my friend Jan is a leftist, since by entertaining that possibility, my awareness already grows in that way.

I think there are two responses to this concern. On the first, we think of updating plans not as commitments that we consciously entertain and adopt, but rather as

⁴Paul herself is sceptical about Diachronic Contenance in general. Michael Bratman (2012) argues in its favour from the assumption that we must, as agents, desire self-governance. I won't pursue the objections to Diachronic Contenance here.

dispositions to update in a particular way. You can easily have a disposition to do something in response to a particular stimulus without being at all aware of the possibility of that stimulus. So the worry from above does not arise if we conceive of updating plans like that. On the second response, we think of updating plans not as mental states that an individual is ever actually in, whether consciously or unconsciously, whether dispositions or intentions, but rather as devices we use to determine the rationality of an individual's posterior. So, we don't imagine that the individual actually has an updating plan and then assess its rationality and demand that they follow through on that plan when their evidence or their awareness grows; instead, we look at how they actually updated and say that they did so rationally if there is any updating plan that it would have been rational for them to adopt had they been able to adopt it, and which would have given rise to their actual posterior had they followed it.

So I think we can make sense of updating plans in the case of awareness growth. But then our extension of Brown's argument runs into a further obstacle. It asks us to use our prior to calculate the expected pragmatic utility of a particular updating plan. To do this, we have to specify the pragmatic utility of an updating plan r at a possible world w in $\mathcal{W}_{\mathcal{F}}$. In the case of evidential growth, we assumed that r is defined on \mathcal{E} , which is a partition of the worlds in $\mathcal{W}_{\mathcal{F}}$. If that's the case, by specifying a world, we thereby specify the evidence we'd learn if that were the actual world, and thereby specify the posterior the updating plan demands at that world. But in the case of awareness growth, if we specify a world w in $\mathcal{W}_{\mathcal{F}}$, we don't thereby specify the way our awareness will grow, because these worlds are only grained as finely as is needed to determine the truth values of the propositions in \mathcal{F} , and the individual does not have opinions about how their awareness will grow, and so the worlds don't specify this. Indeed, that's the moral of the argument above: if they were to have opinions about the ways in which their awareness might grow, their awareness would already have grown in that way. So there is no way to calculate the expected pragmatic utility of an updating plan in the case of awareness growth. As a result, any analogue of Brown's argument places no constraints on your updating plan and therefore no constraints on your posterior. That is, if this argument establishes anything, it's Anything Goes.

Now, you might think that this is too quick.⁵ While the individual does not have credences about how their awareness might grow and so does not have the credences required to calculate the expected pragmatic utility to which Brown's argument appeals, you might think that the credences that they do assign place constraints on any credence function you might use to calculate that expectation on their behalf. For instance, you might think it would be legitimate to calculate that expectation using any credence function that *extends* the individual's prior in such a way that the extension is coherent but now assigns credences to the different possible ways in which their awareness might grow. And in fact, if we do this, Brown's argument will entail a substantial norm that governs how we should respond to awareness growth. It entails Conservatism, which says that you should retain your credences in any proposition to which you assigned credences before your awareness grew. The reason is that, using any such extended credence function in this case is formally indistinguishable from taking this extended credence function to be the individual's prior in Brown's argument for Bayesian Conditionalization and then assuming that the individual is sure to learn a tautological proposition – that is, the partition from which their evidence will come contains just one proposition and it is a tautology. In this case, Brown's argument says that the

⁵Thanks to Alejandro Pérez Carballo for pushing me to consider this proposal.

individual should retain their prior credence function as their posterior one. So *any* credence function that extends the individual's prior will say that this prior will lead them to the best choices in expectation. And that, you might think, gives a good argument for Conservatism.

The problem with this argument is that it isn't clear why only extensions of an individual's prior are legitimate vantage points from which to assess the pragmatic value of some candidate posteriors. In a sense, the very question at issue in this paper asks which vantage points the individual would consider legitimate when they are required to define their credence function over new propositions. So it begs the question in favour of Conservatism to assume that they only legitimate vantage points are those that extend their prior.

4.2. Planning + narrow scope + epistemic argument

The same problem arises for the epistemic version of Brown's argument, which is given by Greaves and Wallace (2006). In Brown's pragmatic argument, we asked which updating plan maximizes expected *pragmatic* utility from the point of view of your prior. In this argument, we ask which maximizes expected *epistemic* utility from that point of view. At a given world, we take the epistemic utility of an updating plan to be the epistemic utility of the posterior it recommends as a response to the evidence you would obtain at that world; and we take the epistemic utility of a credence function at a world to be a measure of how well the posterior does from a purely epistemic point of view at that world. That is, if your epistemic utility function is eu , then the epistemic utility of a credence function c at a world w is $eu(c, w)$, and the epistemic utility of an updating plan r at a world w is $eu(r, w) = eu(r_{E_w}, w)$. These measures of epistemic utility might capture many different epistemic features of the credence function, but they will often take its epistemic utility at a world to be its accuracy at that world, which might be understood as its proximity to the credence function it would be ideal to have at that world, which assigns maximal credence to all truths and minimal credence to all falsehoods.⁶

Greaves and Wallace assume that epistemic utility functions are *strictly proper*. That is, each probabilistic credence function expects every other credence function to have lower epistemic utility than it expects itself to have. If we assume that your measure of epistemic utility is strictly proper, then Greaves and Wallace show:

- (i) any updating plan that is a conditionalizing plan for your prior maximizes expected epistemic utility from the point of view of your prior; and
- (ii) any updating plan that isn't a conditionalizing plan for your prior does not maximize expected epistemic utility from the point of view of your prior.

But now the same problem we faced when we tried to extend Brown's argument to cover cases of awareness growth arises again here. For what is the epistemic utility of an updating plan at a world in this case? In the evidential case, it is the epistemic utility of the posterior that the updating plan demands in response to the evidence that you will learn at that world. But in the awareness growth case, the worlds about which you have an opinion beforehand do not specify the ways in which your awareness might grow,

⁶For details of the accuracy-first view of epistemic value for credences, see Joyce (1998) and Pettigrew (2016).

and so we cannot define the epistemic utility of the updating plan at one of those worlds.

4.3. Planning + wide scope + pragmatic argument

The third argument for Bayesian Conditionalization is due to David Lewis (1999), though I'll present a slight variant. This argument fares better than the previous two. As before, it argues first for planning to condition your prior on your evidence, and secondly for following through on that plan. The first part works by providing an argument for two norms that, together, entail that you should plan to update by conditioning on your evidence. The first norm says that your prior should be a mixture or weighted average of the possible posteriors you might adopt if you follow your updating plan. In symbols:

Weak Reflection Principle (WRP) Suppose

- (i) c defined on \mathcal{F} is your prior credence function;
- (ii) r is your updating plan;
- (iii) r'_1, \dots, r'_n defined on $\mathcal{F}'_1, \dots, \mathcal{F}'_n$ are all the possible credence functions that your updating plan might require you to adopt, where $\mathcal{F} \subseteq \mathcal{F}'_1, \dots, \mathcal{F}'_n$;
- (iv) r_1, \dots, r_n are the restrictions of r'_1, \dots, r'_n to \mathcal{F} .

Then it should be that c is a mixture of r_1, \dots, r_n . That is, there should be non-negative real numbers $\lambda_1, \dots, \lambda_n$ that sum to 1 such that, for all propositions X in \mathcal{F} ,

$$c(X) = \lambda_1 r_1(X) + \dots + \lambda_n r_n(X)$$

The second norm says that you should plan to become certain in whatever evidence you acquire.

Evidential Certainty Suppose r is your updating plan. Then, it should be that, for each E in \mathcal{E} , $r_E(E) = 1$.

It's reasonably straightforward to show that the Weak Reflection Principle and Evidential Certainty entail that your updating plan should be a conditionalizing plan. The second part of the argument is the same as before: we assume something like Diachronic Continence, which we met in section 4.1, and infer Bayesian Conditionalization.

How to establish the Weak Reflection Principle and Evidential Certainty? Here is an argument I have given for the first (Pettigrew 2021):

- (i) if you violate the Weak Reflection Principle, there is a set of bets that your priors will lead you to accept at the earlier time and a set of bets that any of the possible posteriors will lead you to accept at the later time such that, taken together, these will lose you money at all possible worlds relative to your agenda; and
- (ii) if you satisfy the Weak Reflection Principle, there can be no such sets of bets.

And here is the argument for the second:

- (i) if you violate Evidential Certainty, then there are bets you will accept if you learn a certain piece of evidence that will lose you money at all possible worlds relative to your agenda at which you learn that evidence; and
- (ii) if you satisfy Evidential Certainty, there can be no such sets of bets.

So much for the case of evidence growth. How does this argument fare in the case of awareness growth? Well, it's pretty much custom-made for the purpose. Indeed, the Weak Reflection Principle already furnishes us with a norm that governs the updating plans required for this case. However you plan to respond to each of the possible expansions of your agenda, it had better be that your prior is a mixture of the possible posteriors your plan might bequeath to you.

One problem with this argument is that the norm it establishes in the case of awareness growth is slightly toothless. Requiring only that the prior is a mixture of the possible posteriors does nothing to rule out rather bizarre updating plans. Suppose, for instance, that your prior is defined on only two exclusive propositions, *Jan is on the left* and *Jan is on the right*. And suppose that there are two ways your awareness might grow: you might divide *Jan is on the left* into two more fine-grained possibilities, *Jan is a socialist* and *Jan is a communist*, and leave *Jan is on the right* untouched; or you might divide *Jan is on the right* into *Jan is a social conservative* and *Jan is a fiscal conservative* and leave *Jan is on the left* untouched. Suppose you currently assign these credences:

<i>Jan is on the left</i>	<i>Jan is on the right</i>
1/2	1/2

If your awareness growth leads you to divide *Jan is on the left* into *Jan is a socialist* and *Jan is a communist*, you'll adopt these credences:

<i>Jan is a socialist</i>	<i>Jan is a communist</i>	<i>Jan is on the right</i>
3/12	2/12	7/12

If your awareness growth leads you to divide *Jan is a conservative* into *Jan is a social conservative* and *Jan is a fiscal conservative*, you'll adopt these credences:

<i>Jan is on the left</i>	<i>Jan is a social conservative</i>	<i>Jan is a fiscal conservative</i>
7/12	2/12	3/12

Then you satisfy the Weak Reflection Principle. Indeed, in general, that principle permits you to plan to respond to a particular way your awareness might grow in a bizarre and seemingly unmotivated way providing you plan to respond to related ways your awareness might grow in an equally bizarre and unmotivated way, since they then balance out to give your sensible prior!

The problem becomes even more acute if we interpret updating plans as hypothetical posits used to understand the rationality of actual updating behaviour in the way I sketched above. There I said that we might assess the rationality of actual updating behaviour by asking whether there was some rationally permitted updating plan that would lead you to update in this way if you were to follow it. The problem is that, if there is more than one way your awareness might grow, then, for any updating behaviour whatsoever, there is some updating plan defined on the different ways your awareness might grow that recommends that you update exactly as you did if your awareness grows in the way it did. For any prior and any posterior, there is an alternative posterior

such that the prior is a mixture of the two possible posteriors. So, on this interpretation of updating plans in the case of awareness growth, the Weak Reflection Principle imposes no constraints.

4.4. Planning + wide scope + epistemic argument

R. A. Briggs and I have given an epistemic analogue of Lewis' pragmatic argument for Bayesian Conditionalization, and it has recently been corrected and improved by Michael Nielsen (Briggs and Pettigrew 2020; Nielsen 2021). Again, I present a slight variation that goes via my argument for the Weak Reflection Principle. For this, we must assume not only that our measures of epistemic utility are strictly proper as we did above, but also that they are additive. This means that there is, for each proposition X , a measure leu_X of the epistemic utility of assigning different credences to X at different possible worlds, and that the epistemic utility of an entire credence function is the sum of the epistemic utilities of the individual credences it assigns: so, if c is defined on \mathcal{F} , then $\text{eu}(c, w) = \sum_{X \in \mathcal{F}} \text{leu}_X(c(X), w)$. We call leu a *local epistemic utility function* and eu a *global epistemic utility function*. With those assumptions, I obtain the following (Pettigrew 2021):

- (i) if your prior is not a mixture of your possible posteriors, then there is an alternative prior, and, for each possible posterior, an alternative to that, such that, your prior and any of the possible posteriors, taken together, have lower epistemic utility than the alternative prior and the corresponding alternative posterior at any possible world relative to your agenda; and
- (ii) if your prior is a mixture of your possible posteriors, there is no such alternative.

So now we have an epistemic argument for the Weak Reflection Principle. In the case of evidential growth, this can be leveraged into an argument for planning to condition on your evidence and then into an argument for Bayesian Conditionalization. In the case of awareness growth, it already provides a norm. But the concerns I raised about its toothlessness return here.

4.5. Direct + narrow scope + epistemic argument

We turn now to two arguments that try to show directly that you should condition on your evidence, rather than showing first that you should plan to do so and then arguing that you should do what you plan to do. Here again, the pragmatic and epistemic arguments are very similar. I'll begin this time with the epistemic argument, which is due to Dmitri Gallow (2019), improving on an original argument that Hannes Leitgeb and I gave (Leitgeb and Pettigrew 2010). As with the epistemic arguments from the previous section, we begin with a epistemic utility function eu . As before, we assume that it is strictly proper and additive.

Now Gallow thinks that such an epistemic utility function is appropriate if you care about your epistemic utility at all of the possible worlds. But, as our evidence increases, it rules out more and more worlds as possible. And when that happens we should no longer care about the epistemic value of our credences at those worlds. So, for Gallow, our epistemic utility function should change as our evidence changes. At the beginning of our epistemic life, when we have no evidence, it should be strictly proper. But then later, when we have a particular body of evidence, it should match our original epistemic utility function for those worlds at which the evidence is true; but it should

take a constant value of 0 at those worlds at which the evidence is false. By doing that, we encode into our epistemic utility function the fact that we do not care about the epistemic value of our credence function at those worlds that our evidence has ruled out. In symbols: Suppose that, when you have no evidence, your epistemic utility function is eu – that is, $eu(c, w)$ measures the epistemic value of having credence function c at world w . Then, if at some future point your total evidence is given by the proposition E – that is, E is the conjunction of all propositions in your evidence – then your epistemic utility function should be eu_E , which we define as follows:

$$eu_E(c, w) := \begin{cases} eu(c, w) & \text{if } E \text{ is true at } w \\ 0 & \text{if } E \text{ is false at } w \end{cases}$$

Then, Gallow shows that the posterior that maximizes expected epistemic value from the point of view of your prior and when your epistemic utility function is determined by your new evidence in the way just defined is the one demanded by Bayesian Conditionalization.

How might we adapt this argument to apply to the case of awareness growth? Let's take the two types of case, refinement and expansion, in turn. First, refinement. For instance, let's recall the example in which you initially consider the possibilities that Jan is a liberal, a centrist, or a conservative, and then become aware of the distinction within liberalism between ordoliberalism and classical liberalism and so come to divide that possibility in two. Your initial credences are:

<i>Jan is a liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
1/3	1/3	1/3

How should you now set your credences in these four possibilities? Gallow's argument suggests you should maximize expected epistemic utility from the point of view of your priors. So we take each credence function defined on the expanded set of possibilities that includes ordoliberalism, classical liberalism, centrism, and conservatism, and we measure its epistemic utility at each world, and then we weight that epistemic utility by your prior credence in that world, and sum them up to give their expected epistemic utility; and we pick the one with the greatest expected epistemic utility. The problem is that there is a mismatch between the worlds on which your prior credences are defined and the worlds at which the epistemic utility of a credence function on the expanded set of possibilities is defined. The former only specify whether Jan is liberal, centrist, or conservative, while the latter must specify whether she is ordoliberal, classical liberal, centrist, or conservative. So we can use our prior credences only to assess the epistemic utility of posterior credences defined on the original, unexpanded agenda. But we can do that. So what does Gallow's argument tell us about them? Well, since you don't learn any new evidence between the earlier and the later time, according to Gallow, my epistemic utility function should stay the same. And so my prior credences in *Jan is a liberal*, *Jan is a centrist*, and *Jan is a conservative* will expect themselves to be the best posteriors as well. So Gallow's argument seems to suggest that I should assign the same credences to those three possibilities before and after I've come to realize that there are two different ways in which Jan might be a liberal. But of course this is

precisely what I suggested above is not required. This says that there must be some credence p between 0 and $1/3$ such that my posteriors are:

<i>Jan is an ordoliberal</i>	<i>Jan is a classical liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
p	$1/3 - p$	$1/3$	$1/3$

But I suggested above that it would be rationally permissible to assign credence $1/4$ to each. So Gallow's argument seems to prove too much.

Let's consider expansion next. Here, we face a problem. In cases of expansion, there is no set of possibilities on which your priors are defined that we can use to define the expected epistemic utility of the posterior credence functions, even when we restrict those posteriors to the original agenda. After all, what is distinctive about cases of expansion is that you learn that the possibilities that you considered before, and on which your priors are defined, were not exhaustive: in cases of expansion, you expand the set of possibilities considered in your agenda, filling in part of logical space that you hadn't previously considered. So one of the things you learn when your awareness grows by expansion is that any attempt to define an expectation using your priors will fail because you do not have priors defined over a set of possibilities that partitions the logical space.

So, in cases of refinement, Gallow's argument says that we should retain our prior credences in any set of exclusive and exhaustive possibilities. This is the norm that I called Conservatism above. On the other hand, in cases of expansion, it supports nothing stronger than Anything Goes.

4.6. *Non-planning + narrow scope + pragmatic argument*

The pragmatic argument for conditioning on your evidence is identical to Gallow's, except that epistemic utility is replaced by pragmatic utility. As before, we imagine that you will face a decision at the later time after you adopt your updated credence function. In Brown's argument, the pragmatic utility of an updating plan at a world is the pragmatic utility of the posterior credence function it recommends at that world, and the pragmatic utility of a posterior credence function at a world is the utility at that world of the option it leads you to choose. In this argument, the pragmatic utility of a posterior is the same as in Brown's argument at worlds at which your evidence is true; but, like in Gallow's epistemic utility argument, learning evidence rules out possible worlds and leads you no longer to care about the utility of the option you choose at those worlds; so you give every credence function a constant utility of 0 at worlds at which your evidence is false. In symbols: if a^c is the option that maximizes expected utility from the point of view of the credence function c , $u(a^c, w)$ is the utility of that option at world w , and E is your total evidence, then your pragmatic utility function is:

$$u_E(c, w) := \begin{cases} u(a^c, w) & \text{if } E \text{ is true at } w \\ 0 & \text{if } E \text{ is false at } w \end{cases}$$

And then we have:

- (i) for any decision problem you might face, the posterior obtained by conditioning your prior on your evidence maximizes expected pragmatic utility from the point of view of your prior; and

- (ii) for any posterior other than the one obtained by conditioning your prior on your evidence, there is a decision problem you might face for which that posterior does not maximize expected pragmatic utility from the point of view of your prior.

What about the case of awareness growth? What does the argument tell us in that case? Unsurprisingly, it leads to the same conclusions as the epistemic version based on Gallow's argument. In cases of refinement, it supports Conservatism; in the case of expansion, nothing stronger than Anything Goes.

5. Doxastic crises and the normative authority of your prior

So now we have met six arguments for Bayesian Conditionalization. And we have adapted each so that it covers not only the case of evidence growth, but also the case of awareness growth. And we have seen that, when applied to that case, these arguments no longer speak with one voice. The epistemic and pragmatic narrow scope planning arguments of sections 4.1–4.2 place no constraints on your posteriors after your awareness grows. The epistemic and pragmatic wide scope planning arguments of sections 4.3–4.4 impose the Weak Reflection Principle, which places some constraints on how you should plan to update the credences in your old propositions, but constraints weak enough that they don't give rise to constraints on how you should actually update, and which say nothing at all about the credences in the new propositions of which you have become aware. The epistemic and pragmatic narrow scope direct arguments of sections 4.5–4.6 place the strongest constraints on your posterior credences in the propositions in your original agenda. Since only the latter pair of arguments really gives us any chance of establishing a substantial constraint on how to update in the face of awareness growth, I'll focus on that in what follows.

To begin, I'd like to consider an objection to Gallow's epistemic argument for Bayesian Conditionalization and its pragmatic analogue. I think the objection fails, but its failure will be instructive. It goes like this: I start with some prior credences; then I learn some new evidence; but by learning that evidence, I realise that my priors are flawed because they don't take that evidence into account; they don't respect that evidence; therefore, my prior has no normative authority at the later time after I've learned the evidence, and so its expectations have no normative authority at that time, and so I'm not required to pick the posterior at that time that would maximize expected epistemic or pragmatic value from the point of view of my prior. So Gallow's argument fails.

I think this is a poor objection. Here's one response: if you assign non-extremal credences to any proposition, you know that your credence function is flawed in the sense that it is not the ideal credence, which will be 1 if the proposition is true and 0 if it's false. So, when you learn new evidence and thereby see that your prior was flawed, you really learn nothing about its normative authority that you didn't already know. Or, put differently: its normative authority cannot have been based on it being flawless since you always knew it wasn't. But, the objector might reply: in the case you describe, you don't know the specific way in which it is flawed – you don't know anything about how to rectify it, whereas in the cases we're considering, you do. But here's a response to that, offered on behalf of Gallow's argument: What is it about learning the evidence that makes me realise that my prior is flawed? Well, you might think that, when I learn the proposition I do, and I see that my prior

does not assign it maximal credence, I see that my prior is flawed. But why think that, just because I've learned a proposition, I must assign it maximal credence? What is the justification of that norm? Gallow's argument provides an answer to both of these questions. It says that, when you learn the new proposition, you adopt a new epistemic utility function, namely, the one that measures epistemic utility the same way that your old one does for worlds at which the proposition is true, but gives a constant epistemic utility of 0 at worlds at which it is false. And then you note that your prior does not maximize expected epistemic utility from its own point of view when epistemic utility is measured in this new way. And, what's more, it recommends a replacement. It says: I was the right way to go when you valued epistemic utility the way you used to; but now you no longer care about your epistemic utility at certain worlds because your new evidence rules them out, so now I think you should adopt this other credence function instead. Indeed, as Gallow shows, it says you should adopt the credence function obtained from your prior by conditioning on your new evidence. So Gallow's argument tells us why I should think my prior is flawed after I learn the evidence. But it does so on the assumption that my prior retains its normative authority while it is being used to assess possible posteriors using my new epistemic utility function. So the objection fails because it relies on an assumption – namely, when I learn new evidence, I realise my prior is flawed – that itself is best justified by assuming something that the objection denies – namely, when I first learn the evidence and change my epistemic utility function, my prior retains its normative authority to assess the possible posteriors and pick out the one I should adopt.

Nonetheless, the objection raises an important point. In order for the arguments to work, your prior has to retain its normative authority at the later time after you learn the evidence. I think it's wrong to say, as the objection says, that learning new evidence always immediately deprives your prior of its normative authority, but that's not to say that nothing can.

In section 2, we saw that the problem of awareness growth only really arises for a permissivist. So let's suppose permissivism is true. Then, at least for some individuals and some bodies of evidence, the evidence alone does not pick out a unique credence function that rationality requires us to have in response. Let's suppose that I am such an individual with such a body of evidence; and let's suppose I have a particular rational prior in response. So there are other possible priors I might have adopted that would have been rational responses to that evidence. What gives this particular prior normative authority for me? It cannot be that it has any advantage over the other rational ones from the point of view of rationality.⁷ Rather, it must be simply that this is *my* prior; it is the prior *I* adopted from the range of rationally permissible priors. Why does this bestow normative authority on it? Well, because I inhabit these credences; I see the world through them; they record what I think about how the world is. And, so long as they do so, I'm rationally required to use them to make decisions. But I am not rationally required to continue to inhabit them in this way. Things can happen to me that shake me out of these beliefs, things that make me stop inhabiting them; things

⁷Indeed, if we follow Elizabeth Jackson's (2019) sufficientarian approach and think that permissivism is true not because there are many responses to our evidence that are maximally and equally good, but because, while there might be a best response, all responses above a certain level of epistemic goodness are good enough and thereby rational, it could be that my rationally permissible prior is in fact worse than many of the alternative rationally permissible ones.

that make me stand outside them and reconsider them from an external vantage point – not, I should emphasize, some objective vantage point that gives a view from nowhere, since the permissivist tends to deny that there is such a thing, but rather simply a vantage point that doesn't inhabit the beliefs I previously held. Sometimes, I am led to stand outside my beliefs by an unexpected shock to the system. For instance, crises of bodily or mental health, bereavement and the subsequent grief, or political and societal cataclysms can lead us to stand outside the view of the world that we have been inhabiting hitherto, and look down on our beliefs and sometimes abandon them.⁸ Less dramatically, the same can happen when we reflect on ways in which those beliefs were formed in the first place. For instance, we might realise that there is a certain arbitrariness to the credences we adopted at the beginning of our epistemic life and with which we began our epistemic journey. And indeed, in a similar spirit, the same can happen when we reflect on the truth of permissivism itself, if indeed we take it to be a truth. Reflecting on the fact that there are other rationally permissible responses to our evidence might lead us to stand outside our current beliefs and ask whether we wish to retain them. So the normative authority of our prior is conditional on us continuing to inhabit it; but there is no norm that prevents us from no longer inhabiting the credences we have and instead picking others that are also rational.⁹

Now, it seems to me that awareness growth might well precipitate the sort of crisis in belief that leads you to abandon your prior and thus deprive it of its normative authority. After all, the way you set your priors might well have been heavily influenced by the possibilities of which you were aware at the time you set them. Becoming aware of new ones might well make you stand outside the credences at which you've now arrived and decide no longer to follow their lead. And, when it does this, the tweaked version of Gallow's argument will have no force, even in the refinement case.

Note that, while we might know that our awareness will grow in the future, we cannot know in advance the specific way in which it will grow, since to know that is already to have undergone the growth in awareness. So a specific instance of awareness growth will come as an unexpected shock, just like the examples of illness and cataclysm from the previous paragraph, though of course typically less dramatic.

Think of Jan's political affiliations. When you come to realise that there are two different ways in which she might be a liberal, namely, ordoliberalism and classical liberalism, this could well shake you out of your current credences, because it makes you think that, when you set them initially, you were working with a flawed or incomplete conception of the space of possibilities. If this realisation does shake you out of your current credences, then they lose their normative authority, and the fact that you maximize expected epistemic utility from their point of view by retaining your prior credences in the propositions *Jan is a liberal*, *Jan is a centrist*, and *Jan is a conservative* does not entail that you should do that.

I think something similar happens when we are introduced to a sceptical hypothesis, whether it is Descartes' malicious demon hypothesis, or the automaton hypothesis that is intended to induce scepticism about the existence of other minds, or Russell's hypothesis that the world was created only five minutes ago, complete with us and

⁸See, for instance, Havi Carel's argument that bodily illness can force the ill person to step outside many of their previous beliefs and ways of interacting with the world, bringing about what Husserl calls *epoché* (Husserl 1982 [1913]: 20; 1999 [1931]: §32; Carel 2014: section 2).

⁹See Titelbaum (2016) for further discussion of cases in which we change our mind without the catalyst of evidence or awareness growth.

all our apparent memories of times before that. Having never considered the possibility that the external world is an illusion, or that other human bodies do not house minds, or that the world is of an extremely recent vintage and our memories beyond a certain point are not real, I react to becoming aware of it by no longer taking my prior to have normative authority. When Stanley Cavell (1979) talks of the vertigo or terror or anxiety that is induced by your first introduction to a sceptical hypothesis, I think this is partly what he means. The beliefs we have inhabited and which encode our view of the world are called into question wholesale and their normative authority evaporates. Here is Duncan Pritchard (2021: 8) describing a similar phenomenon in his discussion of Cavell:

The metaphor [of vertigo] is apt, for it seems that this anxiety [that Cavell describes] is specifically arising as a result of a kind of philosophical ‘ascent’ to a perspective overlooking our practices, and hence to that extent disengaged from them (as opposed to the ordinary pre-philosophical perspective in which one is unself-consciously embedded within those practices).

In our case, the practices are the prior credences; inhabiting those credences is being unself-consciously embedded within them. Awareness growth can often occasion exactly this sort of philosophical ‘ascent’ to a perspective at which those priors no longer have normative authority.

One other thing that can shake us out of our beliefs is the realisation that they possess a rational flaw. To illustrate how this might happen in expansion cases, consider what happens when you previously considered that Jan might be a liberal, centrist, or conservative, but now realize there’s a fourth possibility, namely, that she’s a leftist. Suppose you assigned these credences at the earlier time:

<i>Jan is a liberal</i>	<i>Jan is a centrist</i>	<i>Jan is a conservative</i>
1/3	1/3	1/3

Now you add to your agenda the proposition that Jan is a leftist. Now, there are (at least) two sorts of betting argument that I can make if I wish to show that your credences are irrational. The most common, as well as the most compelling, is this: we show that your credences will lead you to accept a series of bets that, taken together, will lose you money however the world turns out – that is, they lead you to a sure loss. The less common, and slightly less compelling, is this: we show that your credences will lead you to accept a series of bets that, taken together, will gain you no money however the world turns out, and will lose you money given some ways the world might turn out – that is, they lead you to a possible loss with no possible gain. Now, relative to the original set of possibilities – *Jan is a liberal*, *Jan is a centrist*, *Jan is a conservative* – your credences of 1/3 in each are not vulnerable to a sure loss, and they are not vulnerable to a possible loss with no possible gain. However, relative to the new set of possibilities after the expansion – *Jan is a leftist*, *Jan is a liberal*, *Jan is a centrist*, *Jan is a conservative* – your credences are still not vulnerable to a sure loss, but they are vulnerable to a possible loss with no possible gain. That is, they are vulnerable to the less common sort of betting argument. After all, they will lead you to pay £1 for a bet that pays out £3 if Jan is a liberal and £0 if she’s not; they will lead you to pay £1 for a bet that pays out £3 if she is a centrist and £0 if it not; and they will lead you to pay £1 for a bet that pays out £3 if she’s a

conservative and £0 if not.¹⁰ Now, if she is a liberal, a centrist, or a conservative, these bets, taken together, will cancel out and make you no money, but they will also lose you no money. But if she is a leftist, then they will lose you £3. And this will be true whenever you divide your credences entirely over a set of possibilities that is not exhaustive.¹¹ When you come to realise that the set of possibilities is not exhaustive, you realise that your credences make you vulnerable to such bets, and that should be a catalyst for replacing them.

So there is a number of ways in which awareness growth can precipitate a doxastic crisis that robs your priors of their normative authority. Now, it is also true that new evidence might provoke such a crisis and such a loss of normative authority. And so Gallow's argument does not establish that we should *never* update other than by conditioning our prior on our new evidence; only that we should do that when our priors retain their normative authority after the evidence comes in. Sometimes, if gaining the new evidence leads to a doxastic crisis, we might abandon our prior, pick another that we take to have normative authority, and condition that on our total evidence, knowing that doing so will maximize expected epistemic utility from the point of view of the new prior we've picked. But this will be much rarer than in the case of awareness growth. The reason? We tend not to suffer a doxastic crisis when we learn new evidence because we have typically considered the possibility that we will obtain that specific new evidence in advance of actually obtaining it. On the other hand, while we might consider in the abstract the possibility that we will become aware of further possibilities in the future, we cannot consider specific possibilities of which we might become aware, since by considering them we become aware of them. New possibilities, therefore, take us by surprise and thereby lead us to abandon our priors much more often than new evidence.

It is important to emphasize that doxastic crises and the loss of normative authority that they precipitate are not all or nothing affairs. It is possible that some of our prior credences lose their normative force following a growth in our awareness, while others retain theirs. Indeed, this is by far the most common case. Certainly, if you become aware of a sceptical hypothesis, that might lead you to stop inhabiting any of your prior beliefs; it might lead you to stand outside them all and perhaps to start over from scratch. But in nearly all cases, that won't happen. Rather, the effects will be much more local. For instance, when you become aware of the possibility that Jan is a leftist, you have opinions not just about Jan's political affiliations, but also about what you had for breakfast, about your passwords for your different websites, about your friends' birthdays, their likes and dislikes, about the number of white rhinos left in the wild. When your awareness grows to include the possibility that Jan is a leftist, your prior credences in propositions that concern her political affiliation lose their normative authority, but not your credences concerning what you had for breakfast or the white rhino population. In such a case, Gallow's argument tells us, we should retain these latter credences, since they retain their normative authority and we haven't learned any new evidence and so they expect themselves to be the best posteriors we could have in those propositions. But for our credences concerning Jan's politics, the priors have lost their normative authority, and so we are not rationally required to retain

¹⁰Here, we make the usual assumptions of such arguments, namely, that, for any number S , a credence of p in proposition X will lead you to pay $\mathcal{E}pS$ for a bet that pays out $\mathcal{E}S$ if X is true and $\mathcal{E}0$ if X is false. Or, to be more careful, replace pounds with units of utility.

¹¹This argument is due originally to Abner Shimony (1955). For an introductory presentation, see Pettigrew (2020: 2.7, 3.6).

those. Indeed, sometimes, we will be rationally required to change them – this might happen if, as described above, when we learn of a new possibility, we also learn that our priors are susceptible to a weak betting argument.

When does a growth in awareness precipitate a loss of normative authority for some of your prior credences? On the view I'm sketching, this is a descriptive question, not a normative one. There are norms that govern what posteriors you should adopt if certain credences retain their normative authority; but there are no norms that govern when credences should and shouldn't retain that authority. But I think we can say something descriptive about when this tends to happen.

Roughly speaking, when some of my priors lose their normative authority after awareness growth, it is because I realise that, had I been aware at the beginning of my epistemic life of the possibility of which I've just become aware, I would have adopted credences at that time – sometimes known as my *ur-priors* – that, when conditioned on my evidence to date, would not have resulted in the credences I currently have. This is what happens in Mahtani's example of the singer in the shower, and it's what happens in my example in which I become aware of the possibility that Jan is a leftist. In each case, I assigned equal credence to each of the possibilities of which I was aware. When I then became aware of further possibilities, I realised that, if I'd been aware of them at the outset, I would have divided my credences equally over all of those possibilities, and that would have resulted in different credences at this point. Realizing that, the priors I in fact have lose their normative authority and I am free to adopt new ones, perhaps by setting my *ur-prior* over the new set of possibilities and updating it on my total evidence. On the other hand, when I learn the possibility that Jan might be a leftist, and I consider whether my credences concerning the rhino population would be any different were I to have known of the possibility at the outset, I conclude that they would not.

For this reason, most cases of awareness growth are not very disruptive. They result in some small portion of our credences losing normative authority, but most retain theirs.

6. Accuracy and flipflopping

Before we wrap up, let me consider an objection to the picture of rational credence I've been painting here. According to that account, diachronic norms, such as the requirement to update by conditioning on your evidence, or any requirement to respond to awareness growth in a particular way, is conditional on your prior retaining its normative authority at the later time. And there is no rational requirement to continue to take it to have that normative authority. Many different sorts of event can lead you to stand outside your beliefs and reassess them. Now, suppose there are two credence functions that are rational responses to my current evidence. I have the first at an earlier time. Then, at a later time, having learned nothing new and having not experienced any awareness growth, I come to abandon the normative authority of that first credence function, and I adopt the second one to replace it.¹² According to the picture I've been sketching, there is no irrationality here. And yet the following is true: suppose your epistemic utility function is strictly proper in the way that Greaves and Wallace, Briggs and I, and Gallow assumed. Then there is a third credence function such that, however the world turns out, the total epistemic utility of having this third credence

¹²In the literature on permissivism, this is sometimes known as 'flip-flopping' (Meacham 2014).

function at both the earlier and the later time is greater than the total epistemic utility of having the first credence function at the earlier time and the second credence function at the later time. For instance, suppose I assign a credence only to the proposition *Rain*, which says it will rain tomorrow. Suppose that, at the earlier time, I assign credence p to that proposition; and then, at the later time and after receiving no new evidence, I assign credence q to it instead (where $q \neq p$). Then, for any strictly proper epistemic utility function, there is a credence r that lies between p and q such that, by its lights, I would have been better off assigning credence r to *Rain* at both times. Had I done this, the sum of my epistemic utilities at the earlier and later times would be greater regardless of how the world turns out, that is, regardless of whether it does or does not rain tomorrow.¹³ Surely it follows from this that it is irrational to change your credences between an earlier and a later time? After all, if you do, there is an epistemic life you might have led that is guaranteed to be better than the one you do lead.

I think not. It is true that, if I were in a position to pick my entire epistemic life at the outset, it would be irrational for me to pick the one in which I change credences – where my credence in *Rain* changes from p to q without any new evidence – since there's an alternative that's guaranteed to be better – namely, where I assign credence r at both times. But, having picked credence p at the earlier time, and now sitting at the later time standing outside my belief and asking which credence I should assign at that time, this consideration is irrelevant. After all, I can no longer choose to assign r at both times. I can choose to stick with p , or I can change and adopt q . But sticking with p isn't guaranteed to be better than changing. Suppose q is greater than p . Then, for epistemic utility functions that measure the accuracy of credences, assigning p at both times will be better if it doesn't rain tomorrow, since p is more accurate than q in that situation; but it will be worse if it does rain. So the fact that it would have been better for sure to have r at both times does not tell us that, having chosen p at the earlier time, it's irrational to change to q at the later time.

7. Conclusion

I surveyed six arguments for the most well attested diachronic credal norm, namely, Bayesian Conditionalization, and I asked of each what happens if we try to extend it to cases in which it is not your evidence that grows but your awareness. This resulted in arguments for three norms: Anything Goes, which imposes no constraints on posteriors following cases of awareness growth; the Weak Reflection Principle, which imposes some constraints on updating plans, but little on updating itself; and Conservatism, which places fairly strict constraints on posteriors.

However, I argued that, in fact, these arguments only establish these norms conditionally. They only establish the constraint on your posterior in those cases in which your prior retains normative authority for you. This happens only if you continue to inhabit the view of the world that it encodes, and awareness growth often leads you to abandon that view. I noted that the arguments for Bayesian Conditionalization also establish their conclusion only conditional on the prior retaining its normative authority, but I pointed out that new evidence much less often leads us to stand outside our beliefs and reassess them than awareness growth does. So any norms that follow in the case of awareness growth will apply much less often than those that follow in the case of new evidence.

¹³This is a corollary of the main theorem of Briggs and Pettigrew (2020).

In conclusion, our credal response to awareness growth is less often rigorously constrained than our response to new evidence. While the route I've take to this conclusion is different, the conclusion itself lies close to Steele and Stefánsson's.¹⁴

References

- Bradley R.** (2017). *Decision Theory with a Human Face*. Cambridge: Cambridge University Press.
- Bratman M.** (2012). 'Time, Rationality, and Self-Governance.' *Philosophical Issues (Suppl. Nou's)* 22(1), 73–88.
- Briggs R.A. and Pettigrew R.** (2020). 'An Accuracy-dominance Argument for Conditionalization.' *Nou's* 54 (1), 162–81.
- Brown P.M.** (1976). 'Conditionalization and Expected Utility.' *Philosophy of Science* 43(3), 415–19.
- Canson C. de Ms.** 'The Nature of Awareness Growth.' Unpublished manuscript.
- Carel H.** (2014). 'The Philosophical Role of Illness.' *Metaphilosophy* 45(1), 20–40.
- Cavell S.** (1979). *The Claim of Reason: Wittgenstein, Skepticism, Morality, and Tragedy*. Cambridge, MA: Harvard University Press.
- Gallow J.D.** (2019). 'Learning and Value Change.' *Philosophers' Imprint* 19, 1–22.
- Greaves H. and Wallace D.** (2006). 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility.' *Mind* 115(459), 607–32.
- Husserl E.** (1982) [1913]. *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy: First Book*. The Hague: Martinus Nijhoff.
- Husserl E.** (1999) [1931]. *Cartesian Meditations*. Dordrecht: Kluwer.
- Jackson E.** (2019). 'A Defense of Intrapersonal Belief Permissivism.' *Episteme*. <https://doi.org/10.1017/epi.2019.19>.
- Jaynes E.T.** (2003). *Probability Theory: The Logic of Science*. Cambridge: Cambridge University Press.
- Joyce J.M.** (1998). 'A Nonpragmatic Vindication of Probabilism.' *Philosophy of Science* 65(4), 575–603.
- Karni E. and Vierø M.-L.** (2013). "Reverse Bayesianism": A Choice-Based Theory of Growing Awareness.' *American Economic Review* 103(7), 2790–810.
- Kopec M. and Titelbaum M.G.** (2016). 'The Uniqueness Thesis.' *Philosophy Compass* 11(4), 189–200.
- Leitgeb H. and Pettigrew R.** (2010). 'An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy.' *Philosophy of Science* 77, 236–72.
- Lewis D.** (1999). 'Why Conditionalize?' In *Papers in Metaphysics and Epistemology*, pp. 403–7. Cambridge: Cambridge University Press.
- Mahtani A.** (2021). 'Awareness Growth and Dispositional Attitudes.' *Synthese* 198, 8981–97.
- Meacham C.J.G.** (2014). 'Impermissive Bayesianism.' *Erkenntnis* 79, 1185–217.
- Nielsen M.** (2021). 'Accuracy-Dominance and Conditionalization.' *Philosophical Studies* 178, 3217–36.
- Paris J.B. and Vencovská A.** (1990). 'A Note on the Inevitability of Maximum Entropy.' *International Journal of Approximate Reasoning* 4, 181–223.
- Paul S.K.** (2014). 'Diachronic Incontinence is a Problem in Moral Philosophy.' *Inquiry: An Interdisciplinary Journal of Philosophy* 57(3), 337–55.
- Pettigrew R.** (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettigrew R.** (2020). *Dutch Book Arguments*. Cambridge: Cambridge University Press.
- Pettigrew R.** (2021). 'Bayesian Updating When What You Learn Might Be False.' *Erkenntnis*. <https://doi.org/10.1007/s10670-020-00356-8>.
- Pritchard D.** (2021). 'Cavell and Philosophical Vertigo.' *Journal for the History of Analytic Philosophy* 9(9), 8–22.
- Roussos J.** (2020). *Policymaking under Scientific Uncertainty*. PhD thesis, London School of Economics.
- Shimony A.** (1955). 'Coherence and the Axioms of Confirmation.' *Journal of Symbolic Logic* 20, 1–28.
- Steele K. and Stefánsson H.O.** (2021). *Beyond Uncertainty: Reasoning with Unknown Possibilities*. Cambridge: Cambridge University Press.

¹⁴Many thanks to Conor Mayo-Wilson, Catrin Campbell-Moore, Jason Konek, Joe Roussos, Chloé de Canson, Michele Odisseas Impagnatiello, Jonathan Fiat, Kevin Blackwell, Giacomo Molinari, Shike Zhou, Jack Spencer, Bob Stalnaker, Branden Fitelson, Selina Guter, and a referee for this journal for very helpful comments and questions on previous incarnations of this material.

Titelbaum M.G. (2016). 'Continuing On.' *Canadian Journal of Philosophy* 45(5–6), 670–91.

Williamson T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson J. (2010). *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.

Richard Pettigrew is Professor of Philosophy at the University of Bristol, UK. He completed his PhD in Mathematical Logic there in 2008. His research covers Bayesian epistemology, rational choice theory, and the philosophy of mathematics. He has written five books: *Accuracy and the Laws of Credence* (Oxford University Press 2016), *Who Are Universities For?* (University of Bristol Press 2018), *Choosing for Changing Selves* (Oxford University Press 2019), *Dutch Book Arguments* (Cambridge University Press 2020), and *Epistemic Risks and the Demands of Rationality* (Oxford University Press 2022).

Cite this article: Pettigrew R (2024). How Should your Beliefs Change When your Awareness Grows? *Episteme* 21, 733–757. <https://doi.org/10.1017/epi.2022.33>