

What should the Bayesian do when their awareness grows?

Richard Pettigrew

November 18, 2021

Abstract

Bayesians have a well-developed account of how you should change your credences when you learn new evidence; that is, when your body of evidence grows. What's more, they have a 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 new possibilities; that is, when your awareness grows. In this paper, I consider each of the arguments for the Bayesian's account of how respond to evidence, and I ask whether they can generate a Bayesian account of how to respond to awareness growth. I argue that, with one exception, they can't. I conclude that our credal response to awareness growth is considerably less rigorously constrained than our credal response to new evidence.

The epistemology of credences—or Bayesian epistemology, as it is often known—is concerned with two sorts of norms, synchronic ones and diachronic ones. Synchronic norms govern the credences you have at a given time. Putative examples include: Probabilism, which says that your credences at any time should obey the axioms of the probability calculus; the Principal Principle, which says that your credence in a proposition conditional on its objective chance being a certain probability should just be that probability (Lewis, 1980); and the Principle of Maximum Entropy, which says that, at any time, your credences should have maximal Shannon entropy among all credences that respect your total evidence at that time, where Shannon entropy is a mathematical function taken to measure how unopinionated your credences are (Williamson, 2010). On the other hand, diachronic norms govern the relationships between the credences you have at different times. Putative examples include: Bayesian Conditionalization, which says, if between an earlier and a later time you learn a proposition with certainty and nothing more, then your credences at the later time should be your credences at the earlier time conditional on that proposition; Richard Jeffrey's Probability Kinematics, which tells you how

to respond if your evidence does not come in the form of a proposition learned with certainty, but instead determines the posterior credences you should assign to the propositions in a given partition (Jeffrey, 1965, Chapter 11); and Adam's Conditionalization, which tells you how to respond if your evidence determines not your posterior credences in the elements of a given partition, but your posterior credence in some proposition conditional on another (Wagner, 2003; Bradley, 2005).

Now notice that all of the diachronic norms I just listed tell you how to respond when you receive evidence that comes in different forms. And indeed that is a representative sample of the norms discussed in the literature. But obtaining new evidence is not the only thing that can happen to you that might prompt or justify or demand a change in your credences. You might learn nothing new about the world, but instead acquire a new concept, draw a new distinction, or become aware of a new entity or possibility. For instance, previously, you might have categorised people's political views using the categories *very liberal*, *liberal*, *centrist*, *conservative*, *very conservative*, but now you learn the concept of a *leftist*, and you realise such people do not fall under any of your current categories. Or you might have categorised everyone as on the left-right economic axis and then learned of the libertarian-authoritarian social axis. When one of these things happens to you, you often come to assign credences to propositions you have not considered before: propositions that employ the new concept; propositions that depend on the new distinction or the new possibility or ascribe properties to the newly-discovered entity. Having learned of leftists, I will be able to assign a credence to the proposition that my friend Jan is a leftist; having learned of the social axis of political orientation, I might assign credence to the proposition that Jan is a left-libertarian. It is these sorts of case that I wish to consider in this paper. In the relatively small literature that treats them, they tend to be called cases of *awareness growth*.¹ Note that none of the diachronic norms I listed in the previous paragraph apply to such cases.

Typically, cases of awareness growth divide into two groups: *refinement cases* and *expansion cases*. In a refinement case, the experience leads me to refine a possibility I previously considered. This is what happens, for instance, when I am only aware of the economic axis of political orientation and then become aware of the social axis as well. To present it crudely, I might previously have considered only the two possibilities, *Jan is on the left* and *Jan is on the right*. Learning of the social axis, I now split the first of these possibilities into two and consider *Jan is left-libertarian* and *Jan is left-conservative*; and I split the second in the same way and consider *Jan is right-libertarian* and *Jan is right-conservative*. In an expansion case, on the

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

other hand, I learn of a new possibility that is not a refinement of any possibility I have considered before. This happens when I learn of the political left, having been aware only of liberalism and conservatism before that. Again presenting it crudely, I previously considered only *Jan is liberal*, *Jan is a centrist*, and *Jan is conservative*, but now I add a fourth possibility, namely, *Jan is a leftist*. The question is then: what norms, if any, govern the relationship between the credences I assign before my awareness grows—either by refinement or expansion—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. Knowing different readers favour different levels of formalisation, I will throughout try not to give any mathematical or symbolic presentation of a point without first giving it informally.

2 Impermissivism and awareness growth

Before we get started, it's worth noting that there are certain Bayesians for whom the problem barely arises. First, let's consider *Objective Bayesianism*.² According to proponents of this position, there are really just three fundamental norms that govern your credences, and they are all synchronic. Any diachronic norms are merely consequences of these. The three synchronic norms of objective Bayesianism are those we listed above: Probabilism, the Principal Principle, and the Principle of Maximum Entropy. The idea is this: your total evidence at a time imposes constraints on the credences at that time. The credence function you should have is the one that maximises Shannon entropy among all those credence functions that obey Probabilism and the Principal Principle and which satisfy these evidential constraints. Recall: Shannon entropy is a mathematical measure of how unopinionated a credence function is.

²The version I present here is due to Jon Williamson (2010). Other versions include (Jaynes, 2003) and (Paris & Vencovská, 1990).

So, for instance, you might start your epistemic life assigning credences to just three possibilities, which you take to be exclusive and exhaustive. At this point, you have no evidence and so it imposes no constraints on your credences. Therefore, applying the Principle of Maximum Entropy, you assign credence $\frac{1}{3}$ to each possibility, since that's what maximises Shannon entropy. Then you learn with certainty the proposition that is true only at the first two possibilities. This places on your posterior credence function the constraint that it assign maximal credence 1 to that proposition. Applying the Principle of Maximum Entropy at this later time, but this time including this new constraint imposed on your credences by your evidence, you adopt a posterior that assigns credence $\frac{1}{2}$ to the first two possibilities and credence 0 to the third, since that's what maximises Shannon entropy among the credence functions that satisfy Probabilism, the Principal Principle, and the new constraint.

Now, suppose that, instead of learning that proposition with certainty, you don't acquire any evidence about the world at all, but instead you become aware of a distinction that divides the first possibility into two more fine-grained possibilities, but leaves the second and third possibilities untouched. Then, applying the Principle of Maximum Entropy, you adopt a posterior that divides credence equally across the four possibilities of which you are now aware: the two fine-grained versions of the original first possibility, and then the original second and third possibility. So it assigns credence $\frac{1}{4}$ to each: $\frac{1}{4}$ to each of the fine-grained versions of the first original possibility, and therefore $\frac{1}{2}$ to the first original possibility itself; and $\frac{1}{4}$ to each of the original second and third possibilities.

For the objective Bayesian, there are no fundamentally diachronic norms, and so none that govern cases of awareness growth. The fundamental norms—Probabilism, the Principal Principle, and the Principle of Maximum Entropy—are synchronic, and they are sufficient on their own to determine the unique rational response to your epistemic situation at every time in your epistemic life. These synchronic norms might entail diachronic norms, but the latter are epiphenomenal.

A second sort of Bayesian for whom the problem barely arises is one who thinks there is really just one credal norm, and again it is synchronic. We might call the norm Evidential Probabilism and the position that endorses it *Evidential Bayesianism*. Evidential Probabilism posits an evidential probability function, which takes any body of evidence and any proposition and returns the so-called evidential probability of that proposition relative to that evidence, which is taken to measure how likely that evidence makes that proposition, or how strongly that evidence supports or confirms that proposition (Williamson, 2000). Again, for this sort of Bayesian, any diachronic norms are epiphenomenal—they follow from this single synchronic norm, which determines what credences someone should assign at

a given time without reference to their credences at any other time.

In general, the problem of awareness growth hardly arises for those who are impermissivists about rational credence. An impermissivist says that, for any individual, at any time, and for any proposition to which they assign a credence at that time, there is exactly one credence that rationality requires them to assign to that proposition at that time, and this unique credence does not depend on the individual's previous credences. For more austere impermissivists, this unique credence is determined only by the individual's evidence at that time. The Evidential Probabilist is such an impermissivist. For less austere impermissivists, it is determined partly by the evidence and partly by the set of possibilities the individual considers. The Objective Bayesian is such an impermissivist. Neither sort of impermissivist need give diachronic norms governing awareness growth. Rather, according to them, at each time in your epistemic life, you need not consider what credences you had at any earlier time in order to set your credences now; you need only consider the evidence you have now, and perhaps the set of possibilities you consider; these alone determine the rational credences you should assign to the propositions in your agenda.

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?

One popular such norm is Reverse Bayesian, which says that, while you are not required to retain your credences in the old propositions when you come to assign credences in the new ones, you should at least retain the ratios between any two of the old propositions (Karni & Vierø, 2013).³ In symbols:

Reverse Bayesianism (RB) Suppose c defined on \mathcal{F} is your credence function at the earlier time t , and c' defined on \mathcal{F}' is your credence function at the later time t' , and $\mathcal{F} \subseteq \mathcal{F}'$. And suppose that, between t and t' the only epistemically relevant thing that happens to you is that you become aware of the propositions in \mathcal{F}' . Then, for all X, Y in \mathcal{F} , it should be that

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

Anna Mahtani (ta) has a neat counterexample to RB. You are staying

³Richard Bradley (2017) proposes a closely related principle, which he calls *Rigid Extension*. The counterexamples to Reverse Bayesianism that we will consider below are equally counterexamples to that.

with a friend, Bob, and while alone in the kitchen you hear someone singing in the shower. You assign credence 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 *Tenant*. Knowing Bob is a tenant, you assign credence 1 to $Bob \rightarrow Tenant$. And you assign $\frac{1}{2}$ to *Bob*. 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, whatever credences you assign to *Landlord*, *Tenant*, *Bob*, $Bob \rightarrow Tenant$, *Other*, and $Other \rightarrow Tenant$ (that is, the propositions in \mathcal{F}'), the ratios between your credences in *Landlord*, *Tenant*, *Bob*, and $Bob \rightarrow Tenant$ (that is, the propositions in \mathcal{F}) should be the same as before. But that entails that your new credence in *Bob* must equal your new credence in *Tenant*, since that was the case with your old credences. And thus, if your new credences are to be coherent and your credence in $Other \rightarrow Tenant$ is 1, your new credence in *Other* must be 0. And that doesn't seem right.

There are many examples with the same structure. In each, the situation is akin to the case we described in our presentation of Objective Bayesianism's treatment of awareness growth above: that is, the awareness growth leads the individual to divide certain possibilities they had previously considered into more fine-grained possibilities, but they do not divide each of the original possibilities equally. In Mahtani's case, your realisation that there might be another tenant leads you to divide the possibility *Tenant* in two, but it leaves the possibility *Landlord* untouched. Something similar will happen if you come to realise a logical possibility you hadn't considered before. For instance, you might have originally categorised people's religious beliefs by labelling them *theist*, *agnostic*, or *atheist*. But it might never have occurred to you that anyone might believe in more than one god, so you assign the same credence to *monotheist* and *theist* (and you're certain of $monotheist \rightarrow theist$). When you do realise this alternative possibility, you might divide the possibility *theist* into the more fine-grained

possibilities *monotheist* and *polytheist*, just as you divided *Tenant* into *Bob* and *Other* above.

And something similar will also happen if you enrich your conceptual space in a way that divides some possibilities but not others. For instance, a person who grew up red-green colour blind in a community of red-green colour blind people might only have the colour concepts *yellow*, *blue*, and *green*. So, when they think of some object they haven't seen, and wish to assign credences to its possible different colours, they might assign $\frac{1}{3}$ to *It is yellow*, *It is blue*, and *It is green*. After all, even if we don't think, with the Objective Bayesians that it is rationally required to choose the credences with maximum Shannon entropy, we might nonetheless think it is rationally permitted. But the person might then be told that, among those objects they see as green, there are in fact two colours, green and red; and among those objects they see as yellow, there are in fact two colours, orange and yellow. Or indeed, at some future stage of ocular medicine, they might have their colour blindness surgically corrected. This might lead them to acquire new concepts: they might divide *yellow* into *orange-yellow* and *yellow-yellow*, and *green* into *red-green* and *green-green*, but leave *blue* untouched. And they might then again follow the Objective Bayesian, as they are permitted though not mandated to do, and divide their credences equally over the five new possibilities, namely, *It is orange-yellow*, *It is yellow-yellow*, and so on. But, if they do that, their new credence for *yellow* will be the sum of their new credences for *yellow-yellow* and *orange-yellow*, and that is twice their new credence for *blue*, whereas their old credence in *yellow* is the same as their old credence in *blue*. So they will violate Reverse Bayesianism despite doing something that seems rationally beyond reproach.

For these reasons, I think Reverse Bayesianism must be wrong.

4 Looking to the arguments

When we seek to extend an existing norm to cover a broader range of cases than it currently governs, 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 paper, I'd like to do the latter in an attempt to extend the standard Bayesian norms for diachronic credal norms so that they cover cases in which you respond not to new substantial evidence about the world, but to awareness growth. I'll consider two sorts of arguments for the standard norm of Bayesian Conditionalization, which says that, if between an earlier and a later time, the only epistemi-

cally 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:

Bayesianism Conditionalization (BC) Suppose:

- (i) c is your credence function at the earlier time t , and c is defined on \mathcal{F} ;
- (ii) c' is your credence function at the later time t' , and c' is defined on \mathcal{F} ;
- (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} ,

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

The first sort of argument I'll consider argues not directly for updating by conditioning on your evidence, but for planning to update by conditioning on your evidence. I'll consider both pragmatic and epistemic versions of these arguments (Section 5). The second sort of argument does argue directly for updating by conditioning. Again, I'll consider both pragmatic and epistemic versions (Section 6).

5 Arguments for planning to condition on your evidence

5.1 Pragmatic arguments for planning

Let's begin with the pragmatic arguments for Bayesian Conditionalization. The first is due to Peter M. Brown (1976). As is often pointed out, Brown's argument is not in fact an argument for updating by Bayesian conditioning; rather, it is an argument for planning to update in this way. It goes like this. You must first assume that there is some set of propositions from which the evidence you acquire between an earlier time and a later time will come; and you must assume that this set forms a partition. An *updating plan* is then a function that takes a proposition in this set and returns the posterior credence function the plan asks you to adopt should you learn that proposition with certainty. We say that an updating plan is a *conditionalizing plan*

for your prior if, for each element of the partition to which your prior assigns positive credence, the plan tells you to respond to learning it by conditioning your prior on it. We then imagine that, at the later time, after you've updated your credences in response to the evidence you acquire, you will face a decision between a number of options; and we assume that you will choose between the options by maximising expected utility from the point of view of your credences at that time. We then say that the pragmatic utility of an updating plan at a possible world is the pragmatic utility of the posterior credence function it recommends in response to the evidence you'll receive at that world; and the pragmatic utility of a posterior is the utility of the option that maximises expected utility from its point of view. 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 maximises 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 maximise the expected utility from the point of view of your prior.

The second pragmatic argument for Bayesian Conditionalization is due to David Lewis (1999). Again, it's an argument for planning to update by conditioning, not an argument for simply updating in that way. This time, the argument does not assume you know the partition from which your future evidence will come. Rather, it takes any proposition to which you currently assign a positive credence and asks how you plan to update should you learn this proposition with certainty. And it shows two things:

- (i) if you plan to update in some way other than by conditioning, then there is a set of bets your prior credences will lead you to accept and, whether or not you learn that proposition, there will be a further set of bets that your posterior credences will lead you to accept such that, taken together, the earlier and later sets of bets will lose you money for sure; and
- (ii) if you plan to update by conditioning on the evidence, there can be no such sets of bets.

Can these arguments be adapted to establish any norms that cover the cases of awareness growth we are considering? One problem is that it isn't clear how to formulate the notion of an updating plan in such a case. In Brown's argument, plans are functions defined on partitions that contain the possible pieces of evidence the individual might receive between the

earlier and later times. What should replace the partition in the case of awareness growth? You might naturally think that we should replace that component with the set of possible future agendas to whose members the individual might come to assign credences at the later time. So suppose you now consider the agenda that includes only *Landlord, Bob, Tenant, and Bob* \rightarrow *Tenant*. But at a later point you will expand that either by considering the possibility there is one other tenant and thus adding *Other* and *Other* \rightarrow *Tenant* or by considering that there are another two tenants and thus adding *Other* and *Yet Another*, and *Other* \rightarrow *Tenant* and *Yet Another* \rightarrow *Tenant*. And now suppose we ask how you should plan to respond to these different possible ways in which your awareness might grow. The problem is this: when you consider all of the extra possibilities in these new agendas in order to formulate your updating plans, you already expand your awareness to include them. So you are unable to make these plans without expanding your awareness and setting new credences in the new agendas anyway; you are unable to make these plans without solving the very problem the plans were designed to address. The point is that the different possible ways in which your awareness might expand are very different from the ways in which your evidence might expand. In the latter case, there are some possibilities considered at the earlier time that simply aren't actual and of which you shouldn't become certain—reflecting on possible future evidence shouldn't make you certain of that evidence now. But in the former case, in contrast, by considering the extra possibilities at the earlier time, you immediately expand your agenda to include them.

And the same point applies to Lewis' argument. In that case, our plan isn't defined for all of the different ways we might expand our agenda. Rather, for each one, we make a bespoke plan. But the point still applies. When we make that plan, we become aware of the possibilities of which we might become aware between now and the later time. But we thereby become aware of them now, and so we have to solve the problem of updating in the face of this awareness growth before we can even make the plan for how to do so.

Now, you might complain that I have been too literal in my understanding of updating plans. I have conceived of them as mental items that the individual considers consciously at the earlier time. But perhaps we might better, or at least alternatively, conceive of them as dispositions.⁴ That is, an updating plan says how we are disposed to update our credences when

⁴See Mahtani (ta) for a related appeal to dispositions in the context of awareness growth. In Mahtani's terminology, I am using credence functions to model the credences the individual assigns to propositions of which they're consciously aware; and I'm now considering the dispositions we have to adopt different credence functions upon awareness growth. Mahtani tends to speak instead of dispositions to have different credences in the same propositions in different circumstances. But I think the resulting pictures are very close.

our awareness grows in different ways. If we think of them like that, the argument given in the previous few paragraphs does not apply. I can be disposed to respond to becoming aware of new propositions in a particular way without already being aware of them. And so we might still hope to run Brown's and Lewis' arguments so that they govern the rationality of having certain updating dispositions.

In fact, conceiving of plans as recording dispositions does not help revive Brown's argument. After all, in that argument, you assess the pragmatic utility of an updating plan from the point of view of your prior credences. But that requires that your prior credences assign probabilities to the possible stimuli for your updating—new evidence in Brown's case; in our case, new concepts, distinctions, or hitherto unconsidered logical possibilities, any one of which results in our agenda expanding. But of course it's exactly by having credences in the possible ways in which your awareness might grow that results in awareness growth itself in the way described in the previous few paragraphs. So that's not going to work.

How about Lewis' argument? Here, things seem more hopeful. After all, just as there is a sure loss betting argument for planning to update by Bayesian conditioning, so there is a sure loss betting argument for a norm that I will call the Weak Reflection Principle (van Fraassen, 1999; Pettigrew, *ta*). It says: your current credence function should be a mixture (or weighted sum or convex combination) of your possible future credence functions, once those possible future credence functions have been restricted to the agenda on which your prior is defined. In symbols:

Weak Reflection Principle (WRP) Suppose

- (i) c is your credence function at t , and c is defined on \mathcal{F} ;
- (ii) c'_1, \dots, c'_n are the possible credence functions you might have at t' , and c'_i is defined on \mathcal{F}'_i , where $\mathcal{F} \subseteq \mathcal{F}'_i$;
- (iii) c_i is the restriction of c'_i to \mathcal{F} .

Then it should be that c is a mixture of c_1, \dots, c_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 your original agenda \mathcal{F} ,

$$c(X) = \lambda_1 c_1(X) + \dots + \lambda_n c_n(X)$$

As I note (Pettigrew, *ta*):

- (i) if you violate WRP, 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 however the world turns out; and
- (ii) if you satisfy WRP, there can be no such sets of bets.

So it might seem that this result places normative constraints on how you should be disposed to respond to awareness growth. However you are disposed 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 dispositions might bequeath to you. Nonetheless, I'm a little sceptical. One problem is that it is a little externalist for my taste. That is, you might fall foul of WRP because of some feature of the world of which you are unaware and of which you could not hope to have been aware at the time.

To see this, we look first to the application of WRP in the standard case in which your posterior is a response not to awareness growth but to new evidence. Suppose that, between an earlier and a later time you will learn exactly one of the propositions from a set with certainty. If that set forms a partition, and if, for each proposition in it, the posterior credence function you are disposed to adopt in response to learning that proposition is certain of it, then WRP entails that each possible posterior is obtained from your prior by conditioning on the evidence that prompts it. That is, under this condition on your possible future evidence, WRP entails that you should plan to update by conditioning your prior on that evidence. But now suppose that the set of propositions you might learn between the earlier and later time does not form a partition. Then it does not follow from WRP that you should update by conditioning on your evidence. Indeed, for some such sets of possible future evidence, WRP entails that you should not update in this way! Here's an example: A friend has bought a new brooch, but I don't know what colour it is. I consider three possibilities, *Red*, *Yellow*, *Blue*, and I assign credence $\frac{1}{3}$ to each. Now suppose that, between now and some future time, I will learn something about its colour from another friend who has seen it. But, unbeknownst to me, because that other friend has a particular sort of colourblindness, they will either tell me *Red* \vee *Yellow* or *Red* \vee *Blue*. They will not be able to rule out *Red*, and so they will not tell me *Yellow* \vee *Blue*. And now suppose that I am disposed to react to these two possible pieces of evidence by conditioning on whichever I receive. Then, whichever I receive, my posterior credence in *Red* will become $\frac{1}{2}$. And since my prior was $\frac{1}{3}$, my prior cannot be a mixture of my possible future posteriors. So by being disposed to update by conditioning, I violate WRP and become vulnerable to a sure loss. And the reason is an external constraint on the evidence I might acquire between the earlier and later time. Something about the world of which I'm not aware makes it possible that I will learn the brooch is red or yellow, and possible that I will learn it's red or blue, but not possible that I will learn it's yellow or blue. And because of that, being disposed to update by conditioning on my evidence becomes irrational.

Of course, you might wonder why this is a problem for externalism rather than a problem for BC or for sure loss arguments in general. After all, surely the problem doesn't go away if we imagine that I know, at the earlier

time, that the only possible evidence I might receive is $Red \vee Blue$ and $Red \vee Yellow$. Knowing only that doesn't rule out any of the possibilities, nor indeed tell in favour of any of them over the others. So surely it's still permissible to assign credence $\frac{1}{3}$ to each? That's true, but now what you learn by the later time is no longer just $Red \vee Blue$ and $Red \vee Yellow$, but also the fact that you learned it. So we now have four possibilities instead of three:

- The brooch is red, and I learn it's either red or blue
- The brooch is red, and I learn it's either red or yellow
- The brooch is blue, and I learn it's either red or blue
- The brooch is yellow, and I learn it's either red or yellow

So, if I assign credence $\frac{1}{3}$ to each of the original three possibilities, I must divide my credence of $\frac{1}{3}$ that the brooch is red between the two new, more fine-grained possibilities that make that up. Let's say I split it evenly. Then, if I learn $Red \vee Blue$ and thereby learn that I learn it, my posterior assigns $\frac{1}{3}$ to Red , $\frac{2}{3}$ to $Blue$, and 0 to $Yellow$. And if I learn $Red \vee Yellow$ and thereby learn that I learn it, my posterior assigns $\frac{1}{3}$ to Red , $\frac{2}{3}$ to $Yellow$, and 0 to $Blue$. And this time my prior is a mixture of the two possible future posteriors. Thus, I satisfy WRP and I am not vulnerable to a sure loss.

The same problem haunts us when we apply WRP to the case of awareness growth. Indeed, there are examples very similar in spirit. Suppose that, rather than learning $Red \vee Blue$ or $Red \vee Yellow$ between the earlier and later time, I instead will expand my awareness in one of two ways. I will either become aware of a distinction between two different sorts of reds and two different sorts of yellow, or between two different sorts of red and two different sorts of blue. In each case, I am disposed to respond by dividing credences equally across the five new possibilities—the four new fine-grained possibilities and the one original possibility, which my awareness growth has not further divided. Then, whichever of the two ways my awareness might grow, my posterior will assign credence $\frac{2}{5}$ to Red . But my prior assigns credence $\frac{1}{3}$. So it is not a mixture of my possible posteriors. I thereby violate WRP and I am vulnerable to a sure loss.

That wraps up all I have to say about the pragmatic arguments for planning to condition on your evidence and their implications for responding to awareness growth. What of the epistemic arguments? In the end, these are formally so similar to the pragmatic ones that the same considerations apply.

5.2 Epistemic arguments for planning

In Brown's pragmatic argument, we appealed to measures of the pragmatic utility of an updating plan. At a given world, we took the pragmatic utility of an updating plan to be the pragmatic utility at that world of the posterior it recommends as a response to the evidence you'll obtain at that world; and we took the pragmatic utility of a credence function at a world to be the utility at that world of the option we would choose from a given decision problem if we were to choose by maximising expected utility from the point of view of that credence function.

In the epistemic arguments, in contrast, we appeal to measures of the epistemic utility of an updating plan. 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. 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 gradational accuracy or proximity to the truth at that world.⁵ That is, they will often measure how close the credence function lies to the epistemically ideal credence function at the world in question; that is, how close it lies to that world's omniscient credence function, which assigns maximal credence to all truths and minimal credence to all falsehoods.

We will consider two properties of epistemic utility functions. First, *strict propriety*. An epistemic utility function is *strictly proper* if each probabilistic credence function expects every other credence function to have lower epistemic utility than it expects itself to have. Second, *additivity*. An epistemic utility function is *additive* if the epistemic utility it assigns to a credence function is the sum of the epistemic utilities of the individual credences that the credence function assigns.

Now, the analogue of Brown's pragmatic argument is Greaves and Wallace's (2006) epistemic argument. Suppose 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 maximises 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 maximise expected epistemic utility from the point of view of your prior.

⁵For details of the accuracy-first view of epistemic value for credences, see (Joyce, 1998) and (Pettigrew, 2016a).

The analogue of Lewis' pragmatic argument is Briggs and my (2020) epistemic argument, which has recently been corrected and improved by Michael Nielsen (ta). Suppose your measure of epistemic utility is strictly proper and additive. Then Briggs and I (and Nielsen) show:

- (i) if you plan to update other than by conditioning, there is an alternative prior and updating plan that, taken together, have greater epistemic utility in all possible worlds than your prior and updating plan have, taken together; and
- (ii) if you plan to update by conditioning on the evidence, there is no such alternative.

The problem we face when we try to extend Greaves and Wallace's as well as Briggs and my (and Nielsen's) epistemic arguments to the case of awareness growth is exactly the same as the problem we face when we try to extend Brown's or Lewis' pragmatic argument. In order to plan how to respond to awareness growth, you must consider the ways in which your agenda might expand in the future; but once you've considered that, it should expand in all of those ways right now. What's more, just as there is a sure loss argument for the Weak Reflection Principle, so there is an epistemic argument (Pettigrew, ta). Suppose your measure of epistemic utility is strictly proper and additive. Then I show:

- (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; and
- (ii) if your prior is a mixture of your possible posteriors, there is no such alternative.

But the concerns I raised about the move to WRP above return here.

So, in the end, I think the arguments for planning to condition on your evidence tell us little about how we should plan to respond to awareness growth, or how we should be disposed to respond.

6 Arguments for conditioning on your evidence

6.1 The epistemic argument for conditioning

So we turn now to arguments that try to show not (only) that you should plan to condition on your evidence, but that you should in fact condition on it. 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 by me and Hannes Leitgeb (Leitgeb & Pettigrew, 2010). As with the epistemic arguments from the previous section, we begin with an epistemic utility function, which takes a credence function and a possible world and returns a measure of the epistemic value of having that credence function at that world. And, as before, we assume that it is strictly proper. So every probabilistic credence function expects itself to do best. 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 \mathbf{EU} —that is, $\mathbf{EU}(c, w)$ 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 , then your epistemic utility function should be \mathbf{EU}_E , which we define as follows:

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

Then, Gallow shows that the posterior that maximises 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. In symbols: If \mathbf{EU} is strictly proper and $c(E) > 0$, then $c(-|E)$ maximises expected epistemic utility from the point of view of c and when your epistemic utility function is \mathbf{EU}_E .

6.2 The pragmatic argument for conditioning

The pragmatic argument for conditioning on your evidence is identical, 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 maximises expected utility from the point of view of the credence function c , $a^c(w)$ is the utility of that option at world w , and E is your total evidence, then your pragmatic utility function is:

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

And then we have:

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

6.3 Adapting the arguments to the case of awareness growth

Now, how might we adapt these arguments 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 Mahtani’s example from above. You begin with credence $\frac{1}{2}$ in each of *Landlord* and *Tenant*, credence $\frac{1}{2}$ in *Bob*, and credence 1 in *Bob* \rightarrow *Tenant*. Then you come to realise that there might be another tenant, so you expand your agenda to include *Other* and *Other* \rightarrow *Tenant*. How should you now set your credences in the original four propositions as well as these two new ones? Gallow’s argument suggests you should maximise expected epistemic utility from the point of view of your prior. But how should we define that expectation? In order to define an expectation, we need a set of possibilities. These possibilities should be pairwise exclusive and collectively exhaustive, so that there is no way in which two of them might both be true, and no way in which all of them might be false; we should assign credences to each; and those credences should sum to 1. Given this, it seems our possibilities must be *Landlord* and *Tenant*, since *Landlord*, *Bob*, and *Tenant* are not exclusive, and precisely what our awareness growth has taught us is that *Landlord* and *Bob* are not exhaustive. Now, suppose we try to calculate the expected epistemic utility of different possible posteriors over our expanded agenda

using that set of possibilities. Then we hit a roadblock. The problem is that, in order to determine the epistemic value of a credence function that assigns credences to *Bob* and *Other*, a possibility must determine the truth value of those two propositions; but, while *Landlord* determines that they are both false, *Tenant* does not determine which of them is true. So it seems that, at best, we might use these possibilities to define the expected epistemic utility of a credence function defined only on *Landlord* and *Tenant*, and not to define the expected utility of a credence function defined on the full new agenda. But even this restricted application gives some substantial results. After all, since I don't learn any new evidence between the earlier and the later time, my epistemic utility function will stay the same. And so my prior credences in *Landlord* and *Tenant* will expect themselves to be best. So Gallow's argument seems to suggest that I should assign the same credences to *Landlord* and *Tenant* before and after I've come to realise that there is another way in which *Tenant* might be true. But of course this is precisely what I suggested above is not required. Realising that there is a third possible singer in the shower, I might naturally divide my credences equally over the three possibilities, thereby making *Tenant* twice as likely as *Landlord*. So Gallow's argument seems to prove too much. Below, we'll see why that might be.

Let's consider expansion next. In this case, we hit a roadblock much sooner. The problem is that, in cases of expansion, there is no set of possibilities that we can use to define the expected epistemic utility of the posterior credence functions, even when we restrict those posteriors to smaller agendas. After all, what is distinctive about cases of expansion is that you learn that the possibilities that you considered before 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 it seems at first sight that, in cases of refinement, Gallow's argument says that we should retain our credences in any set of exclusive and exhaustive possibilities to which we assign credences that sum to unity; whereas in cases of expansion, it is silent and supports no norm. In fact, I think things are a bit more complicated. To see why, let's think about a natural objection to Gallow's argument for updating by conditioning on your evidence.

6.4 Doxastic crises and the normative authority of your prior

You might think: I start with a prior credence function; then I learn some new evidence; but by learning that evidence, I realise that my prior is flawed because it doesn't take that evidence into account; 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 can't be required to pick the posterior at that time that would maximise expected epistemic value from the point of view of my prior.

I think this is a poor objection. To see why, ask this: 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 maximise 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, and 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 Gallow's argument to work, your prior has got 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.⁶ Rather, it must be simply that this is *my* prior; it is the prior *I* picked from the range of possible rational 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 that make me stand outside them and reconsider them from an external vantage point. Sometimes, this happens because of unexpected shocks to the system. For instance, crises of mental health, 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. 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.⁷

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. Think of the case of the singing in Bob's bathroom. When you come to realise that there might be another tenant in the flat, this could well shake

⁶Indeed, 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.

⁷See (Titelbaum, 2016) for further discussion of cases in which we change our mind without the catalyst of evidence or awareness growth.

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 maximise expected epistemic utility from their point of view by retaining your prior credences in *Landlord* and *Tenant* 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 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, let's recast Mahtani's example as such a case. That is, at the earlier time, you have credences only in *Landlord* and *Bob*. You do not consider the propositions *Tenant* or *Bob* \rightarrow *Tenant*. And you assign equal credence of $\frac{1}{2}$ to these two possibilities. As a result, when you come to realise that there might be someone in the house other than Bob or his landlord, you do not refine the possibility *Tenant*, since you do not consider that proposition, but rather you add a new possibility you hadn't considered

before, namely, *Other*. 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—*Landlord* and *Bob*—your credences of $\frac{1}{2}$ 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—*Landlord*, *Bob*, and *Other*—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 £5 for a bet that pays out £10 if *Landlord* is true and £0 if it is false; and they will lead you to pay £5 for a bet that pays out £10 if *Bob* is true and £0 if it is false.⁸ Now, if *Landlord* is true, these bets, taken together, will cancel out and make you no money, but they will also lose you no money; and similarly if *Bob* is true. But if *Other* is true, then they will lose you £10. And this will be true whenever you divide your credences entirely over a set of possibilities that is not exhaustive.⁹ 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 maximise 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. And the reason is related to the points I made above about why we cannot hope for a pragmatic or

⁸Here, 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 $\mathcal{L}pS$ for a bet that pays out $\mathcal{L}S$ if X is true and £0 if X is false. Or, to be more careful, replace pounds with units of utility.

⁹This argument is due originally to Abner Shimony (1955). For an introductory presentation, see (Pettigrew, 2020, 2.7, 3.6).

epistemic argument for planning to update in a particular way in response to awareness growth. 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.

7 Responding to a doxastic crisis

In the previous section, we saw how a doxastic crisis might rob your prior of its normative authority. When it does, I argued, you are free to pick a new prior and then bring it up to date it by maximising your expected epistemic utility from its point of view, where your epistemic utility is defined using your current evidence in the way that Gallow suggests—that is, you pick a new prior and condition it on your current evidence. Are there any constraints on the new prior you should pick? Of course, you must pick a rational one, and which boast that property will depend on your favoured brand of permissivism. But are there constraints beyond that?

Here is one view on which there is. According to a view I've developed elsewhere, the priors you should pick at the beginning of your epistemic life should be determined partly or wholly by considerations of epistemic risk (Pettigrew, 2016b,c, 2022). The idea is this: according to epistemic utility theory, what it is epistemically rational to believe is what you would choose to believe, if you were able to choose what to believe, and if, when choosing, you were moved entirely by considerations of the epistemic value of your beliefs. That is, we do not assume that we are able to choose our beliefs; rather, we assess their rationality by assessing what it would be rational to choose were we able to, and were we to care only about epistemic value when we made the choice. This is the account of epistemic rationality behind Gallow's argument for updating by conditioning on your evidence. But it might also tell us which priors are rational for an individual. In Gallow's argument, when you have to choose your posteriors, you use your priors to calculate the expected epistemic utility of the different possible options. When you're picking your prior, you don't have such a prior available. So you need a decision rule that, unlike maximising expected utility, does not require probabilities as one of its ingredients. Now, there are plenty of these around, and many of them determine what it is rational to choose by appealing to your attitudes to risk. For instance, Abraham Wald's (1945) Minimax decision rule is suited to the maximally

risk-averse: it tells you to choose an option whose worst case outcome is best; that is, you rank the options by the utility each achieves at the worlds where it achieves its lowest utility, and you pick one at the top of the ranking. And, in previous work, I showed that, if you apply this rule when you pick your prior, and if your measure of epistemic utility is strictly proper in the sense defined above, then you should pick the uniform distribution over the set of possible worlds of which you are aware—that is, if there are n such worlds, each should receive a credence of $\frac{1}{n}$ (Pettigrew, 2016b). Another example: noting that Wald’s Minimax principle is too severely risk averse, Leonid Hurwicz (1951, 1952) proposed his Criterion of Realism, which generalises Wald’s rule: it does not demand that you consider only the worst case outcomes of an option, but rather asks that you set a weight for best cases and a weight for worst cases, rank the options by the weighted sum of their best- and worst-case utilities, and pick an option from the top of the ranking. This decision rule accommodates both the more risk averse, who will give greater weight to the worst case, as well as the more risk-inclined, who will give greater weight to the best case—and Wald’s rule is the special case in which you give all of your weight to the worst case. And I have shown that, when we use this rule to pick our priors, it gives different recommendations depending on the weight you give to the best case (Pettigrew, 2016c). Suppose λ is the weight you give to the best case, and $1 - \lambda$ the weight you give to the worst; and suppose n is the number of possible worlds of which you are aware. Then, if $\lambda \leq \frac{1}{n}$, then Hurwicz’s Criterion agrees with Minimax: you should pick the uniform distribution over the worlds, which assigns $\frac{1}{n}$ to each. But, if $\frac{1}{n} < \lambda$, then you should do something different: first, pick a world; second, assign credence λ to that world; third, then divide the remaining credence equally among the remaining worlds, giving $\frac{1-\lambda}{n-1}$ to each.

Now, just as we might use Minimax or Hurwicz’s Criterion at the beginning of our epistemic lives, when we pick our initial credences, so we might use it after a doxastic crisis, when our previous credence function no longer has normative authority over our choice of posterior. And it seems that the sort of event that leads us to abandon the normative authority of our prior is unlikely to lead us to abandon our attitudes to epistemic risk. If that’s so, then those attitudes to epistemic risk might continue to place constraints on the new credence function we pick after our doxastic crisis to replace our abandoned prior, just as they placed constraints on our initial choice of credence function.

Let’s briefly sketch an example. I begin with credences over just two possibilities: *Red*, which says that my friend’s brooch is red, and *Not Red*, which says it isn’t. And perhaps I am maximally risk averse, and so adhere to Wald’s Minimax. Then I should assign credence $\frac{1}{2}$ to each possibility. But now suppose that, without gaining any new evidence, I become aware of two fine-grainings of the possibility *Not Red*: they are *Blue* and *Yellow*.

This leads me to abandon the normative authority of my prior, but not my commitment to extreme risk aversion in the epistemic sphere. So now Minimax tells me to assign credence $\frac{1}{3}$ to *Red*, $\frac{1}{3}$ to *Blue*, and $\frac{1}{3}$ to *Yellow*. And of course, as we noted above, that leads me to violate Reverse Bayesianism, since my initial credence in *Red* is the same as my initial credence in *Not Red*, while my new credence in *Red* is half my new credence in *Not Red*.

Another example, briefly. Again, I am initially aware only of *Red* and *Not Red*. But I am not so risk averse as to use Minimax. Perhaps instead I assign a weight of $\lambda = \frac{5}{12}$ to the best case and $1 - \lambda = \frac{7}{12}$ to the worst case. Then, since $\frac{5}{12} < \frac{1}{2}$, Hurwicz's criterion says that I should assign credence $\frac{1}{2}$ to each possibility. But now suppose again that, without gaining any new evidence, I become aware of two fine-grainings of *Not Red*, namely, *Blue* and *Yellow*. Then, since $\frac{1}{3} < \frac{5}{12}$, Hurwicz's criterion tells me to pick a possibility, assign credence $\frac{5}{12}$ to that possibility, and divide the remaining credence equally between the other two. So I might assign $\frac{5}{12}$ to *Blue*, for instance, and $\frac{7}{24}$ to *Red* and $\frac{7}{24}$ to *Yellow*.

8 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.¹⁰ According to the picture I've been sketching, there is no irrationality here. And yet the following is true: suppose you measure epistemic value using a strictly proper scoring rule of the sort described above. Then there is a third credence function such that, however the world turns out, the total epistemic utility of having this third credence 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

¹⁰In the literature on permissivism, this is sometimes known as 'flip-flopping' (Meacham, 2014).

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. Then, for any strictly proper epistemic utility functions, there is a credence r that lies between p and q such that 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. 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. 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.

9 Conclusion

Arguments for diachronic credal norms come in three varieties. The first, exemplified by Lewis' sure loss argument and Briggs and Pettigrew's (and Nielsen's) epistemic utility argument, take an individual's prior credences and their plans for their posterior credences together, and ask from an impartial standpoint whether there is a problem with having both. The problem we face if we try to extend this to the case of awareness growth is that it isn't clear how we can plan to update in the face of future growth in our agenda without expanding our agenda in that the time we do the planning, which makes the planning futile. The second sort of argument, exemplified by Brown's pragmatic argument and Greaves and Wallace's epistemic argument, takes the priors and the plans for the posteriors and asks the priors to evaluate the plans for the posteriors. This faces the same problem. The third sort of argument, exemplified by Gallow's epistemic argument and

the pragmatic version I described above, does not involve plans. Rather, it takes the prior to evaluate posteriors using an epistemic utility function that is determined partly by the individual's evidence. Extending this argument to the case of pure awareness growth, where there is change in agenda but no simultaneous change in evidence, we get an argument for a very strong conclusion in the case of refinement, where it says that you should retain your credences in the propositions in any partition to which you originally assigned credences; but, in the case of expansion, it supports no norm because there is no partition to which you assigned credences at the earlier time, and therefore no way to calculate expected epistemic or pragmatic utility.

However, I argued that, even in those cases in which this sort of argument applies and seems to establish a norm that governs awareness growth, it only does so conditionally; it only establishes that norm if your prior retains normative authority for you, which happens only if you continue to inhabit the view of the world that it encodes. I noted that this is also true in ordinary cases in which you learn new evidence but do not change your agenda. 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 such norms that follow in the case of awareness growth will apply much less often than those that follow in the case of new evidence.

References

- Bradley, R. (2005). Radical Probabilism and Bayesian Conditioning. *Philosophy of Science*, 72, 342–364.
- Bradley, R. (2017). *Decision Theory with a Human Face*. Cambridge, UK: Cambridge University Press.
- Briggs, R. A., & Pettigrew, R. (2020). An accuracy-dominance argument for conditionalization. *Noûs*, 54(1), 162–181.
- Brown, P. M. (1976). Conditionalization and expected utility. *Philosophy of Science*, 43(3), 415–419.
- Canson, C. (ms). The Nature of Awareness Growth. Unpublished manuscript.
- Cavell, S. (1979). *The Claim of Reason: Wittgenstein, Skepticism, Morality, and Tragedy*. Cambridge, Mass.: Harvard University Press.
- Gallow, J. D. (2019). Learning and Value Change. *Philosophers' Imprint*, 19, 1–22.

- Greaves, H., & Wallace, D. (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459), 607–632.
- Hurwicz, L. (1951). A Class of Criteria for Decision-Making Under Ignorance. Tech. Rep. Statistics 356, Cowles Commission Discussion Paper.
- Hurwicz, L. (1952). A criterion for decision making under uncertainty. Tech. Rep. Technical Report 355, Cowles Commission.
- Jackson, E. (2019). A Defense of Intrapersonal Belief Permissivism. *Episteme*, (pp. 1–15).
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge, UK: Cambridge University Press.
- Jeffrey, R. (1965). *The Logic of Decision*. New York: McGraw-Hill.
- Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4), 575–603.
- Karni, E., & Vierø, M.-L. (2013). “Reverse Bayesianism”: A Choice-Based Theory of Growing Awareness. *American Economic Review*, 103(7), 2790–2810.
- Leitgeb, H., & Pettigrew, R. (2010). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science*, 77, 236–272.
- Lewis, D. (1980). A Subjectivist’s Guide to Objective Chance. In R. C. Jeffrey (Ed.) *Studies in Inductive Logic and Probability*, vol. II. Berkeley: University of California Press.
- Lewis, D. (1999). Why Conditionalize? In *Papers in Metaphysics and Epistemology*, (pp. 403–407). Cambridge, UK: Cambridge University Press.
- Mahtani, A. (ta). Awareness growth and dispositional attitudes. *Synthese*.
- Meacham, C. J. G. (2014). Impermissive Bayesianism. *Erkenntnis*, 79, 1185–1217.
- Nielsen, M. (ta). Accuracy-dominance and conditionalization. *Philosophical Studies*, (pp. 1–20).
- Paris, J. B., & Vencovská, A. (1990). A note on the inevitability of maximum entropy. *International Journal of Approximate Reasoning*, 4, 181–223.
- Pettigrew, R. (2016a). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

- Pettigrew, R. (2016b). Accuracy, Risk, and the Principle of Indifference. *Philosophy and Phenomenological Research*, 92(1), 35–59.
- Pettigrew, R. (2016c). Jamesian epistemology formalised: an explication of ‘The Will to Believe’. *Episteme*, 13(3), 253–268.
- Pettigrew, R. (2020). *Dutch Book Arguments*. Elements in Decision Theory and Philosophy. Cambridge, UK: Cambridge University Press.
- Pettigrew, R. (2022). *Epistemic Risk and the Demands of Rationality*. Oxford, UK: Oxford University Press.
- Pettigrew, R. (ta). Bayesian updating when what you learn might be false. *Erkenntnis*.
- 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*. Ph.D. thesis, London School of Economics.
- Shimony, A. (1955). Coherence and the Axioms of Confirmation. *Journal of Symbolic Logic*, 20, 1–28.
- Steele, K., & Stefánsson, H. O. (2021). *Beyond Uncertainty*. Cambridge Elements in Decision Theory and Philosophy. Cambridge University Press.
- Titelbaum, M. G. (2016). Continuing on. *Canadian Journal of Philosophy*, 45(5-6), 670–91.
- van Fraassen, B. C. (1999). Conditionalization, A New Argument For. *Topoi*, 18(2), 93–96.
- Wagner, C. G. (2003). Commuting probability revisions: The uniformity rule. *Erkenntnis*, 59, 349–64.
- Wald, A. (1945). Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 46(2), 265–280.
- Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.