Conditionalization Does Not (in general) Maximize Expected Accuracy

MIRIAM SCHOENFIELD
The University of Texas at Austin / New York University
mschoenfield@austin.utexas.edu

Abstract: Greaves and Wallace argue that conditionalization maximizes expected accuracy. In this paper I show that their result only applies to a restricted range of cases. I then show that the update procedure that maximizes expected accuracy in general is one in which, upon learning P, we conditionalize, not on P, but on the proposition *that we learned P*. After proving this result, I provide further generalizations and show that much of the accuracy-first epistemology program is committed to KK-like iteration principles and to the existence of a class of propositions that rational agents will be certain of if and only if they are true.

1. Introduction

Rational agents revise their beliefs in light of new information they receive. But how should agents revise their beliefs in response to new information? To state this question more precisely, it will be helpful to think of information processing as occurring in two (not necessarily temporal) stages:[1] First, there is a non-inferential stage at which an agent, through some non-inferential means, gains some information. We'll call this *exogenous* information gaining. Metaphorically, we can think of this stage as involving the world 'flinging' some information at the agent.

In the second stage, the agent revises her beliefs in response to the exogenous information gaining (the flinging) that took place. These are the revisions that we are interested in evaluating. Sometimes, as a result of such revisions, the agent may come to possess additional information, in which case we'll say that this information came to be possessed *endogenously*. For example, I may gain the information that Gabe is at the party exogenously, and, as a result of revising my beliefs in response to this information, also come to (endogenously) possess the information that his partner Henry is at the party.

[1] The two-stage model is discussed (or implicit) in much of the literature on Bayesian updating. See, for example, Howson and Urbach (1989, p.285), Jeffrey (1992, p.38), Bronfman (2014, p.872) and Miller (forthcoming).

More precisely, then, the question we're interested in is this: how does an ideally rational agent revise her opinions in light of the information she receives exogenously*?*

According to Bayesian epistemology, rational agents[2] revise their credences by conditionalization. Informally, conditionalizing on E involves setting your new credence in every propositions, P, to what your old credence in P was on the supposition that E. Formally, you conditionalize on E if

$$p_{new}(\cdot) = p_{\text{old}}(\cdot \mid E)$$

where

$$p(A|B) = p(A\&B) \,/\, p(B).$$

Since conditionalizing is an operation performed on a proposition, thinking of conditionalizing as a way of responding to new information requires characterizing each possible body of information an agent might receive as a proposition. Since one of the aims of this paper is to evaluate an argument for the claim that conditionalizing is the rational response to gaining information, I will assume for now (as is standard) that any body of information that an agent receives exogenously can be uniquely characterized as a proposition (one that is often a conjunction of many other propositions).[3] Later we'll see what happens if we relax this assumption.

The proposition that uniquely characterizes the entire body of information the agent exogenously receives is sometimes referred to in the literature as 'the strongest proposition one learns'. To emphasize the exogenous aspect, however, I will sometimes call this proposition 'the strongest proposition one exogenously learns'. For short, I will sometimes just call it 'the proposition one exogenously learns' or 'the proposition one learns'.

Note that what I am taking as primitive is the notion of exogenously gaining information. I am using the term 'the strongest proposition one exogenously learns'

---

[2] Unless stated otherwise, when I talk about rational agents, I mean ideally rational agents. I discuss non-ideal agents in §4.

[3] Why uniquely? Because if there were more than one proposition that characterized the body of information the agent receives, then the claim that one should conditionalize on *the* proposition that characterizes one's new information wouldn't make sense. If one claimed that one should conditionalize on *a* proposition characterizing this information, then conditionalization would no longer output a unique credence function given an agent's priors and the new information she received. Conditionalization, then, would no longer count as an update procedure in the sense that is necessary for the arguments under discussion.

as a technical term, which presupposes that any body of information can be uniquely characterized as the sort of thing (a proposition) that one can conditionalize on.

Conditionalization is the process of revising one's credences by conditionalizing on the strongest proposition one exogenously learns. Why think that conditionalization is a rational way of revising one's credences? There are a variety of arguments that have been offered,[4] but the focus of this paper will be an argument by Hilary Greaves and David Wallace (2006) for the claim that conditionalization maximizes expected accuracy.

The Greaves and Wallace argument is part of a larger philosophical program that has been of increasing interest: accuracy-first epistemology. The basic tenet of accuracy-first epistemology is that accuracy is the fundamental epistemic value, and the central project that accuracy-firsters pursue involves the derivation of rational requirements from accuracy based considerations.[5] A cluster of accuracy based arguments for rational requirements, including arguments for the requirement to conditionalize, rely on the following claim:

> RATACC: The rational update procedures are those that maximize expected accuracy according to a strictly proper scoring rule.

(The terms used in this principle will be defined precisely in what follows).

I will argue that Greaves and Wallace's result applies only to a restricted range of cases. Thus, even if RATACC is true, Greaves and Wallace's argument does not show that, in general, conditionalizing on the proposition one learns is the update procedure that is rational.

So the question then arises: which update procedure maximizes expected accuracy in general? I show that, in fact, what maximizes expected accuracy in general is not conditionalization, but a rule that I will call 'conditionalization*'. Conditionalization* has us conditionalize on the proposition *that we learn P*, when P is the proposition we learn.[6] I will show that conditionalization* happens to coincide

---

[4] See, for example, Teller (1976), Williams (1980) and van Fraassen (1989, p.331-7) and (1999).

[5] For an overview, see Pettigrew (2016).

[6] I borrow the term 'conditionalization*' from Hutchison (1999). Hutchison describes a class of cases that have been thought to pose problems for conditionalization. One proposal he describes (though does not commit to) for how to deal with these cases is to deny that conditionalization is the rational update procedure. Rather, he proposes, perhaps what's rational, upon learning P, is conditionalizing on the proposition *that we learn P*. Defenders of conditionalization have offered alternative ways of

with conditionalization in the special cases that Greaves and Wallace consider, but it yields different results in all other cases. So my central thesis is the following:

> *Central Thesis*: If RatAcc is true, then the rational update procedure is conditionalization*, and not conditionalization.

I will not, in this paper, evaluate the merits of RatAcc or the accuracy-first program. This is why my central thesis is a conditional claim.

After arguing for this thesis, I discuss some of the interesting implications of my results for iteration principles in epistemology. In particular, I show that if RatAcc is true, it follows that, if we learn P, we're rationally required to be certain that we learned P. I then show that, regardless of how we think about exogenously gaining information, it follows from RatAcc that there is a class of propositions that rational agents will be certain of if and only if they are true. Since many of the results of the accuracy-first program rely on RatAcc, those who deny these claims cannot accept much of what accuracy-first epistemology has to offer.

2. Setup

What does it mean to say that an update procedure maximizes expected accuracy? In this section I lay out the formal framework that I will use to prove the main result.

2.1 Accuracy and expected accuracy

Accuracy is measured by a scoring rule, **A,** which takes a state of the world, s, from a partition of states, S, and a credence function *c* defined over S, from the set of such credence functions, C_S, and maps the credence function/state pair to a number between 0 and 1 that represents how accurate the credence function is in that state.

$$\mathbf{A}: C_S \times S \rightarrow [0,1]$$

Intuitively, we can think of the accuracy of some credence function as its 'closeness to the truth'. *c* is maximally accurate if it assigns 1 to all truths and 0 to all falsehoods. It is minimally accurate if it assigns 1 to all falsehoods and 0 to all truths.

---

treating the cases that Hutchison describes, though Hutchison raises worries for these proposals. My paper provides an independent argument for Hutchison's proposal that doesn't appeal to the controversial cases discussed in his paper.

If an agent does not know which state obtains she will not able to calculate the accuracy of a credence function *c*. However, if she is probabilistically coherent, she will be able to calculate the *expected* accuracy of *c.* (Throughout, I will be assuming that rational agents are probabilistically coherent). The **expected accuracy** of credence function $c \in C_S$ relative to a probability function $p \in C_S$ is:

$$EA^p(c) = \sum_{s \in S} p(s)\, \mathbf{A}(c, s)$$

That is, the expected accuracy of a credence function *c* relative to *p* is the average of the accuracy scores *c* would get in the different states that might obtain, weighted by the probability that *p* assigns to those states obtaining.

A **strictly proper scoring rule** is a scoring rule with the feature that every probability function maximizes expected accuracy relative to itself. In other words, if **A** is strictly proper, then the quantity:

$$EA^p(c) = \sum_{s \in S} p(s)\, \mathbf{A}(c, s)$$

is maximized when $c = p$. I will not argue here for the claim that our accuracy measures should be strictly proper. I will simply assume this to be true in what follows since the accuracy based argument for the claim that we should conditionalize (in addition to other arguments in accuracy-first epistemology[7]) requires strict propriety.[8] See Greaves and Wallace (2006), Gibbard (2008), Joyce (2009), Moss (2011), Horowitz (2013) and Pettigrew (2016) for a discussion of the motivation for using strictly proper scoring rules.

---

[7] For example, the argument for probabilism. See Pettigrew (2016).

[8] Although the accuracy based argument for the claim that conditionalization is the rational update procedure requires strict propriety, it's worth noting that Greaves and Wallace state their main result slightly more generally: rather than assuming RATACC and that the scoring rule is strictly proper, they remain neutral on propriety and assume that the rational update procedures will be those in which one adopts a credence function that is *recommended* by a credence function yielded by an update procedure that maximizes expected accuracy. As a result, their main argument does not show that conditionalization is always rational, but rather, that what they call *quasi-conditionalization* is always rational. In their Corollary 2, they point out that that if we assume that the scoring rule is strictly proper, conditionalization always maximizes expected accuracy, and so is always rational. It is also true that if we assume that the scoring rule is strictly proper, their constraint on rational update procedures is equivalent to RATACC. In this paper, I'm interested in arguments for the claim that conditionalizing (rather than quasi-conditionalizing) is always rationally required and, for these purposes, RATACC and strict propriety must be assumed.

2.2 Learning experiences and update procedures

We're trying to figure out how to revise our credences in light of the exogenous information we gain. What exactly is involved in gaining information? Greaves and Wallace don't say much about this, and I too will remain as neutral as possible. All that is being assumed (by Greaves and Wallace and myself) is that the body of information one exogenously receives can be uniquely characterized as a proposition.

Suppose you know that you're going to undergo some experience, E. E might be waking up tomorrow, or arriving at the office. Assuming you are probabilistic, for any proposition P, the set {P, ~P} is a partition of your possibility space. (A **partition** of a probabilistic agent's possibility space is a set of propositions that the agent regards as mutually exclusive and jointly exhaustive.) So the following is a partition of your possibility space: {I gain some new information upon undergoing E, I don't gain any new information upon undergoing E}. We can represent this partition as follows:

I gain some new information upon undergoing E.	I don't gain new information upon undergoing E.

Now consider all of the possibilities in which you gain new information upon undergoing E. Call these bodies of information: i_1, i_2… You can further subdivide the region in which you gain new information as follows:

I gain i_1	I gain i_2	I gain i_3	I gain i_4	…	I don't gain new information upon undergoing E.

Since we are assuming for now that we can uniquely characterize each possible body of information that you gain as a proposition, and we are describing the possibility in which you gain a body of information as a case in which you *learn* that proposition, we can redescribe the partition above as follows:

I learn $\mathcal{X}_1$	I learn $\mathcal{X}_2$	I learn $\mathcal{X}_3$	I learn $\mathcal{X}_4$	…	I don't gain new information upon undergoing E.

(Recall that 'I learn $\mathcal{X}_i$' is short for: $\mathcal{X}_i$ is the strongest proposition I exogenously learn.)

We'll let L(P) name the proposition that P is the strongest proposition you exogenously learn upon undergoing E. For ease of notation, we'll describe the possibility in which you gain no new information as a case in which you learn the tautology ($\mathcal{T}$). So yet another redescription of the partition above is:

L($\mathcal{X}_1$)	L($\mathcal{X}_2$)	L($\mathcal{X}_3$)	L($\mathcal{X}_4$)	…	L($\mathcal{T}$)

We'll call an event in which an agent exogenously learns a proposition **a learning experience** (and note that, given our terminology, it is consistent with this that the agent 'learns' the tautology and so gains no new information)**.** Now suppose that an agent is considering some learning experience that she will undergo. She can represent her future learning experience by the set of propositions that she assigns non-zero credence to exogenously learning. So we'll say that an agent whose possibility space is as depicted above represents her future learning experience by the set:

$$\mathcal{X}: \{\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}_3 \ldots, \mathcal{T}\}$$

I will sometimes use the name of the set that represents an agent's learning experience as a name for the learning experience itself.

It will be useful for what follows to note that, in general, if $\mathcal{X}$ represents an agent's future learning experience, and L($\mathcal{X}$) is the set of propositions L($\mathcal{X}_i$) for each $\mathcal{X}_i \in \mathcal{X}$, then L($\mathcal{X}$) is a partition of the agent's possibility space.

Here's why: First, imagine a case in which the agent is certain that she will gain some new information upon undergoing the learning experience. Then she will be certain that there will be exactly one proposition in $\mathcal{X}$ that uniquely characterizes the new information that she will exogenously receive. Thus, she will be certain that exactly one member of L($\mathcal{X}$) is true. So if the agent is certain that she will gain some new information, L($\mathcal{X}$) is a partition of her possibility space. If, on the other hand, the agent leaves open the possibility of gaining no new information, then $\mathcal{T}$ will be a member of $\mathcal{X}$. Since our agent is certain that she will gain no new information (learn $\mathcal{T}$) or gain some new information (learn exactly one of the $\mathcal{X}_i$ that is not $\mathcal{T}$), but not both, she too is certain that exactly one proposition in L($\mathcal{X}$) is true. Thus, whether the agent leaves open the possibility of gaining no new information or not, L($\mathcal{X}$) is a partition of the agent's possibility space.

An **update procedure,** *U*, in response to a learning experience, $\mathcal{X}$, is a function that assigns a probability distribution to each member of $\mathcal{X}$, with the intended interpretation that an agent conforming to *U* adopts $U(\mathcal{X}_i)$ as her credence function if and only if the proposition she learns upon undergoing the learning experience is $\mathcal{X}_i$. In other words, on the intended interpretation, an agent conforming to *U* adopts $U(\mathcal{X}_i)$ if and only if L($\mathcal{X}_i$) is true. The fact that an update procedure is a mapping from the propositions the agent might learn to probability functions guarantees that update procedures satisfy what Greaves and Wallace call 'availability': In any two worlds in which the agent learns the same information, the update procedure recommends the same credence function. Conceiving of update procedures in this way is motivated by the thought that what an agent is rationally required to do in response to learning a proposition must be determined completely by which proposition she learns. Later in the paper we'll consider generalizations of the notion that don't take this assumption for granted.

It will sometimes be convenient to think of *U* as assigning to each possible *state* a credence function. So we can let *U*(s) be $U(\mathcal{X}_i)$, where $\mathcal{X}_i$ is the proposition that the agent learns in state s.

$$U(s) = U(\mathcal{X}_i) \text{ where } s \in L(\mathcal{X}_i)$$

As we'll see in a moment, what Greaves and Wallace call 'an experiment' is just a special kind of learning experience, and what Greaves and Wallace call 'an available act' is just an update procedure in response to an experiment. So my notions are generalizations of the notions that Greaves and Wallace use.

2.3 Experiments and available acts

Greaves and Wallace's discussion assumes that the agent contemplating her future learning experience satisfies the following two conditions:

> PARTITIONALITY: The propositions that the agent assigns non-zero credence to exogenously learning form a partition of the agent's possibility space.
>
> FACTIVITY: The agent is certain that if she learns P, P is true. [9]

In cases in which PARTITIONALITY and FACTIVITY hold, we will say that the agent's future learning experience is representable as an **experiment**.

Greaves and Wallace's definition of an **available epistemic act *A*** is: 'an assignment of a probability distribution to each piece of possible information $\mathcal{E}_j \in \mathcal{E}$ [where $\mathcal{E}$ is a partition] with the intended interpretation that if $A(\mathcal{E}_j) = p_j$ then p_j is the probability function that an agent performing act *a* would adopt as his credence distribution if he received the new information that the actual state was some member of $\mathcal{E}_j$' (p. 611-612). Thus, an available act is just an update procedure in response to an experiment.

Now, if *every* rational agent satisfied PARTITIONALITY and FACTIVITY, then perhaps it wouldn't matter that Greaves and Wallace's result only applies to such agents (for their account could still be a general account of how to revise *rational* credence functions). So it's worth thinking about whether a rational agent may fail to satisfy these conditions.

To begin, note that, *prima facie*, it would be quite surprising if all rational agents satisfied PARTITIONALITY. To return to our flinging analogy, imagine that the world has a 'bucket' of propositions $\{\mathcal{X}_1, \mathcal{X}_2 \ldots\}$ that you think it might fling at you. If you know that the world will fling exactly one proposition in the bucket at you, then the set: {the world flings $\mathcal{X}_1$, the world flings $\mathcal{X}_2$, the world flings $\mathcal{X}_3$…} is, indeed, a partition of your possibility space. But so far we've been given no reason to think that the propositions *in the bucket itself* form a partition of your possibility space. After all, what if the bucket contains both P and P&Q? Since P&Q entails P, any set that contains P&Q and P is not a partition. This means that if an agent leaves open the possibility that P is the strongest proposition she exogenously learns, and also leaves open the possibility that P&Q is the strongest proposition she exogenously

[9] Greaves and Wallace are explicit about PARTITIONALITY, but not FACTIVITY. However, as we'll see, FACTIVITY must be assumed for their arguments to work.

learns, then the agent doesn't satisfy PARTITIONALITY. But it's hard to see why it would be irrational for an agent to leave open the possibility that the strongest proposition she learns is P, and also leave open the possibility that the strongest proposition she learns is P&Q.

To illustrate the strength of the claim that all rational agents satisfy PARTITIONALITY and FACTIVITY, it will be helpful to prove the following lemma (I call it a 'lemma' because it will play an important role in a proof that comes later):

> Lemma 1
>
> An agent satisfies PARTITIONALITY and FACTIVITY if and only if, for each $\mathcal{X}_i$ such that she assigns non-zero credence to $\mathcal{X}_i$ being the strongest proposition she exogenously learns, the agent assigns credence 1 to:
>
> $$L(\mathcal{X}_i) \leftrightarrow \mathcal{X}_i$$
>
> Proof
>
> Suppose that PARTITIONALITY and FACTIVITY are satisfied. FACTIVITY entails that the agent assigns credence 1 to the left-to-right direction of the biconditional: $L(\mathcal{X}_i) \rightarrow \mathcal{X}_i$ for any $\mathcal{X}_i$. For FACTIVITY says that, for all $\mathcal{X}_i$, the agent is certain that if she learns $\mathcal{X}_i$, $\mathcal{X}_i$ is true. What about the right-to-left direction? If PARTITIONALITY holds, then the agent is certain that exactly one proposition in $\mathcal{X}$ is true. Since, by assumption, the agent is certain that she will learn one proposition in $\mathcal{X}$, and that (due to FACTIVITY) it will be a true proposition, she will have to learn *the* one true proposition in $\mathcal{X}$. So if $\mathcal{X}$ forms a partition, she is certain that the $\mathcal{X}_i$ that is true is the proposition that she will learn. This gives us: $\mathcal{X}_i \rightarrow L(\mathcal{X}_i)$. Thus, any agent that satisfies PARTITIONALITY and FACTIVITY will, for each $\mathcal{X}_i \in \mathcal{X}$, assign credence 1 to $L(\mathcal{X}_i) \leftrightarrow \mathcal{X}_i$.
>
> Conversely, suppose that for every proposition $\mathcal{X}_i$ that an agent assigns non-zero credence to learning, she assigns credence 1 to: $L(\mathcal{X}_i) \leftrightarrow \mathcal{X}_i$. And recall that the $L(\mathcal{X}_i)$ form a partition of the agent's possibility space.[10] It follows

[10] See §2.2 for the detailed argument for this, but here's the gist: $L(\mathcal{X}_i)$ is the proposition that the strongest proposition an agent exogenously learns is $\mathcal{X}_i$. So an agent can't leave open the following possibility: For distinct $\mathcal{X}_1$ and $\mathcal{X}_2$, the strongest proposition I exogenously learn is $\mathcal{X}_1$ and the strongest proposition I exogenously learn is $\mathcal{X}_2$. This is because, assuming $\mathcal{X}_1$ and $\mathcal{X}_2$ are distinct, if the agent exogenously learns $\mathcal{X}_2$, then it's false that the strongest proposition she exogenously learns is $\mathcal{X}_1$. Since the agent can't leave open the possibility that there are two propositions that are each

that an agent who regards the $\mathcal{X}_i$ as equivalent to the $L(\mathcal{X}_i)$ will be such that the $\mathcal{X}_i$ *also* form a partition of the agent's possibility space. So any agent who is certain that, for each $\mathcal{X}_i$, $L(\mathcal{X}_i) \leftrightarrow \mathcal{X}_i$, satisfies PARTITIONALITY. And under the assumption that the agent is certain that, for each $\mathcal{X}_i$, $L(\mathcal{X}_i) \rightarrow \mathcal{X}_i$ (which is the just the left-to-right direction of the biconditional), it follows that the agent satisfies FACTIVITY as well: she is certain that if she learns some proposition, $\mathcal{X}_i$, that proposition is true. Thus, any agent that is certain that, for each $\mathcal{X}_i$, $L(\mathcal{X}_i) \leftrightarrow \mathcal{X}_i$ satisfies PARTITIONALITY and FACTIVITY.

So the question of whether a rational agent could fail to satisfy PARTITIONALITY or FACTIVITY amounts to the following question: Might there be some proposition, P, such that a rational agent assigns non-zero credence to exogenously learning P, while leaving open the possibility that P is true, though she won't learn it, OR leaving open the possibility that she will learn P, but P isn't true?

Let's begin by considering the first type of case: a case in which an agent leaves open the possibility that P, but she doesn't learn that P.

<u>P but not L(P)</u>

Seemingly, there are many cases in which, for some P that I might learn, I leave open the possibility that P is true though I don't learn it. Suppose, for example, that I am about to turn on my radio and am considering the possible bodies of information I might receive. I think that one possibility is that I learn:

R: It is raining in Singapore.

and nothing else. I also think, however, that it might be raining in Singapore even if I don't learn that it is when I turn on the radio. This seems perfectly rational, but if so, then it is rational to leave open the possibility that R (a proposition I might learn) is true but I don't learn that it is.

In response, one might claim that it is, in fact, irrational for me to leave open the possibility that I exogenously learn R and nothing else. For perhaps one thinks

---

the strongest proposition she exogenously learns, the agent must think that at most one member of the set $L(\mathcal{X}_i)$ is true. She will also think that at least one member of the set is true since we are assuming that she is certain that she will undergo a learning experience represented by $\mathcal{X}$: that is, she is certain that she will learn one member of $\mathcal{X}$. (Recall that this is consistent with her leaving open the possibility of gaining no new information and merely 'learning' the tautology.) Thus, she will be certain that at least one member of $L(\mathcal{X})$ is true and that at most one member of $L(\mathcal{X})$ is true.

that I should be certain that any case in which I come to exogenously possess the information that R as a result of turning on the radio is a case in which the strongest proposition that I exogenously learn is something like:

> R(R): It is being reported on the radio that it is raining in Singapore.

So, one might claim, if I am certain that I will turn on the radio, I should be certain that if R(R) is true, I will learn that it is.

But should I? What if I leave open the possibility that upon turning on the radio all I will hear is static? In that case I might leave open the possibility that it is being reported on the radio that it is raining in Singapore, even if I don't learn that it is being reported on the radio that it is raining in Singapore. Surely it is not irrational to leave such a possibility open.

In response to this, one might claim that it is also irrational for me to think of R(R) as a proposition in the bucket of propositions that the world might fling at me (that is, as a potential strongest proposition I exogenously learn). Rather, one might claim, the proposition in the vicinity that I should assign non-zero credence to exogenously learning is:

> E(R(R)): I have an experience as of it being reported on the radio that it is raining in Singapore.

And perhaps, one thinks, I am rationally required to be certain that if E(R(R)) is true, I will learn that it is.

Note, however, that for this this strategy to generalize the following two claims must be true:

(a) If P is a proposition about one's experience (that one could, in principle, learn about), then a rational agent should regard it as impossible for P to be true without her learning that P.

(b) Every agent should assign credence zero to P being the strongest proposition she exogenously learns, unless P is a proposition about her own experience.

Why is (b) necessary? Because it's plausible that for any proposition P that is *not* about an agent's experiences, an agent can rationally leave open the possibility that P is true though the agent doesn't learn that it is. So if agents are to be certain that *all* propositions they might learn will be true *only if they learn them*, they must be certain that the only kinds of propositions they will exogenously learn are propositions about their experience. Why is (a) necessary? Because claiming that the only propositions I learn are about my experience will be of no help if I can leave open the possibility that some proposition about my experience is true but I don't learn that it is.

But (a) and (b) are far from obvious. Let's begin with (a). Consider, for example, the following proposition:

> Detailed-E(R(R)): I have an experience as of a reporter with a British accent saying that it is raining in Singapore with a slight emphasis on the word 'raining' and a pause between 'raining' and 'Singapore'.

This seems like a proposition I could learn. But it also seems possible that my experience could have the described features and yet I don't exogenously learn that it does. I may not *notice* the accent, or the pauses, or the emphases, despite the fact that these features are present in my experience. So why couldn't a rational agent leave open the possibility that a proposition like this is true, though she doesn't learn that it is?

(b) is also a very substantive assumption. Why should every agent be antecedently certain that propositions about her experience are the *only* kinds of propositions she will exogenously learn? Presumably small children exogenously learn things: the world flings bodies of information at them. But small children might not even have the conceptual apparatus that makes it possible for them to exogenously learn propositions about their own experience. So one might want to claim that children, at least, can exogenously learn propositions that are not of this sort. But if the world can fling propositions like R, or R(R), into a child's belief box, what should make me antecedently *certain* that the world won't fling such a proposition at me? In other words, if propositions that aren't about one's experience can, in principle, be exogenously learned, why should every agent be certain that she won't undergo this sort of learning?

In sum, while there is nothing incoherent about the view that, for any proposition P one might learn, one is rationally required to be certain that if P is true, one will learn it, such a view requires some rather hefty commitments about the kinds of propositions that can be exogenously learned. The resulting commitments are stronger than even the kinds of luminosity commitments that (some) internalists are happy to sign up for and that Timothy Williamson (2000) and others have argued against. For it's not just that one can't be wrong about one's own experiences. And it's not just that, for some class of experiences, having the experience always puts one *in a position* to know that one is having it. It's not even that, whenever some proposition is true of one's experience, one *in fact* comes to know that proposition. It is that *every* rational agent must *antecedently* be *certain* that *any* proposition P that could be true of her experience (and which it is possible to learn about) is a proposition that she *will* learn exogenously whenever P is true and that there are no other propositions that she could exogenously learn.

<u>L(P) but not P</u>

If you think that the word 'learn' is factive, and that any rational agent should be certain of this, you might think that a rational agent can never leave open the possibility of learning a proposition that is false. But let's set aside the semantics of 'learn'. For various reasons, some philosophers have thought that an agent might have a false proposition as part of her *evidence*.[11] So if we redescribed the project as an investigation into how an agent should revise her credences in light of the *evidence* she receives (instead of 'in light of what she exogenously *learns*'), we might want an account that allows a rational agent to leave open the possibility of gaining a false proposition as part of her evidence. In this case, we would want an account that would apply to agents that fail to satisfy FACTIVITY.

Given the considerations above, I think that it should remain a live possibility that a rational agent may fail to satisfy one of PARTITIONALITY or FACTIVITY. So if we want a fully general account of credal revision, we should consider how such agents should revise their credences in light of what they learn. This forces us to consider learning experiences that aren't representable as experiments.

2.4 The expected accuracy of update procedures

[11] See, for example, Rizzierie (2011), Arnold (2013), Comesaña and McGrath (2014) and Drake (forthcoming).

So far, we have defined the expected accuracy of a credence function. But we don't yet have a definition of 'the expected accuracy of an update procedure in response to a future learning experience'. Greaves and Wallace do provide such a definition. However, Greaves and Wallace's definition can only be used to describe the expected accuracy of an update procedure for an agent that satisfies PARTITIONALITY and FACTIVITY. Since, in this paper, I am interested in which update procedures maximize expected accuracy in general, I will have to generalize their notion.

So what *do* we mean by the expected accuracy of an update procedure U in response to a future learning experience $\mathcal{X}$? On an intuitive level, what we're trying to capture is how accurate we expect to be upon learning a member of $\mathcal{X}$ if we conform to U. And recall that, on the intended interpretation, an agent conforms to U if she adopts $U(\mathcal{X}_i)$ whenever the proposition she learns upon undergoing the learning experience is $\mathcal{X}_i$.

Suppose that an agent knows that she will undergo a learning experience represented by $\mathcal{X}$. Let $\mathbf{A}(U(s),s)$ represent the accuracy score of an agent conforming to U in s. It is natural to think of the *expected* accuracy that an agent assigns to U as the weighted average of the accuracy scores that an agent conforming to U would adopt in each state in which she learns a member of $\mathcal{X}$. This gives us the following understanding of the expected accuracy of an update procedure: The **expected accuracy of an update procedure *U*** in response to a future learning experience $\mathcal{X}$, relative to an agent's probability function p is:[12]

$$EA^p(U) = \sum_{s \in L(\mathcal{X})} p(s)\, \mathbf{A}(U(s), s)$$

$$= \sum_{L(\mathcal{X}_i) \in L(\mathcal{X})} \sum_{s \in L(\mathcal{X}_i)} p(s) * \mathbf{A}(U(\mathcal{X}_i)), s)$$

[12] My definition of expected accuracy is inspired by the definition provided by Greaves and Wallace (though there is one important difference, the reason for which will become clear shortly). A limitation of defining expected accuracy using summations is that if the number of things being summed over is infinite, the sum may not be defined. Kenny Easwaran (2013) provides an alternative way of understanding the notion of expected accuracy that coincides with Greaves and Wallace's definition when finite quantities are involved, but also applies to cases when the quantities are infinite. The results that follow can be carried out in Easwaran's framework (see note 15). However, since the crucial points in this paper are most easily brought out using the Greaves and Wallace inspired definition, I will continue using summations in the main text.

I will now prove a second lemma:

Lemma 2

If an agent's future learning experience is representable as an experiment, $\mathcal{E}$, and U is an update procedure in response to $\mathcal{E}$, then:

$$EA^{p}(U) = \sum_{L(\mathcal{E}_i)\in L(\mathcal{E})} \sum_{s\in L(\mathcal{E}_i)} p(s)^{*}\, \mathbf{A}(U(\mathcal{E}_i)), s) = \sum_{\mathcal{E}_i\in\mathcal{E}} \sum_{s\in\mathcal{E}_i} p(s)^{*}\, \mathbf{A}(U(\mathcal{E}_i)), s)$$

Proof

Note that the first (leftmost) double sum is just the definition of the expected accuracy of an update procedure. The second double sum is just like the first except that, rather than summing over the $L(\mathcal{E}_i)$, we're summing over the $\mathcal{E}_i$.

We know from Lemma 1 that if an agent's future learning experience is representable as an experiment – that is, the agent satisfies PARTITIONALITY and FACTIVITY – then the agent is certain that for all propositions $\mathcal{E}_i \in \mathcal{E}$:

$$\mathcal{E}_i \leftrightarrow L(\mathcal{E}_i)$$

Given this, there is no harm in replacing the '$L(\mathcal{E}_i)$' that features in the definition of the expected accuracy of an update procedure with '$\mathcal{E}_i$'.

Since Greaves and Wallace assume PARTITIONALITY and FACTIVITY, they can simply *define* the expected accuracy of an update procedure (which they call 'an act') in response to an experiment as the average accuracy scores that would result from adopting $U(\mathcal{E}_i)$ whenever $\mathcal{E}_i$ *is true*. And this, indeed, is what they do. Their definition of the expected accuracy of an act corresponds to the double sum on the right-hand side of the lemma. But it's important to realize that they *wouldn't* define expected accuracy this way if they weren't assuming PARTITIONALITY and FACTIVITY. This is because, without these assumptions, the double sum on the right does not represent a weighted average of the scores that would result from an agent performing the act. For Greaves and Wallace, in defining an act, say that an agent performs act U in response to $\mathcal{X}$ if she adopts $U(\mathcal{X}_i)$ as her credence function if and only if *she learns* $\mathcal{X}_i$ (p.612). But if an agent leaves open the possibility that $\mathcal{X}_i$ is true, though she doesn't learn it (PARTITIONALITY fails), or that she learns it, though it's not true (FACTIVITY fails), then an agent performing *U* would *not* adopt $U(\mathcal{X}_i)$ if and only if $\mathcal{X}_i$ is true. Thus, it is only if PARTITIONALITY and FACTIVITY are assumed that the double sum on the right represents the expected accuracy of the credences that result from an agent performing *U*.

2.5 Summing up

The purpose of this section was to develop a precise definition of the notion of the expected accuracy of an update procedure in response to a learning experience. Although Greaves and Wallace provide a definition for the expected accuracy of an *act* in response to an *experiment*, this definition won't apply to cases in which PARTITIONALITY or FACTIVITY fail.

I defined the expected accuracy of an update procedure as the weighted average of the accuracy scores that would result form an agent conforming to the update procedure (adopting $U(\mathcal{X}_i)$ whenever she learns $\mathcal{X}_i$). I then showed that if the agent can represent her future learning experience as an experiment, this quantity will equal the weighted average of the accuracy scores that would result from her adopting $U(\mathcal{X}_i)$ whenever $\mathcal{X}_i$ is *true*. This gives us Greaves and Wallace's definition of the expected accuracy of an act. Thus, my framework, in terms of update procedures and learning experiences, is a generalization of the framework developed by Greaves and Wallace.

In the next section I will use the generalized framework to derive Greaves and Wallace's result: the claim that, *for an agent who can represent her future learning experience as an experiment*, conditionalizing on the proposition she learns maximizes expected accuracy. I then prove a more general result: for *any* agent contemplating a future learning experience, the update procedure that maximizes expected accuracy is one in which, upon learning $\mathcal{X}_i$, the agent conditionalizes on the proposition *that she learned* $\mathcal{X}_i$. In cases in which the learning experience is representable as an experiment (and only in such cases), this amounts to the same thing as conditionalizing on $\mathcal{X}_i$.

3. The Greaves and Wallace Result and its Generalization

Greaves and Wallace argue that (given a strictly proper scoring rule) conditionalizing on the proposition one learns is the update procedure that maximizes expected accuracy in response to an experiment.

We can think of the argument for this claim as involving two steps. First, there is a purely formal result that demonstrates that plugging in certain values in certain quantities maximizes those quantities. Second, there is an argument from this formal result to the claim that, given our understanding of update procedures, expected accuracy of update procedures, learning, and experiments, the update procedure (or available act) that maximizes expected accuracy in response to an experiment is the one that has the agent conditionalize on the proposition she learns. It will be important to keep these two steps separate. I will call the purely formal result that can be extracted from Greaves and Wallace's paper 'G&W'.

> G&W: For any partition of states P: $\{P_1...P_n\}$, consider the set of functions, $\mathcal{F}$, that assign members of P to probability functions. The member of $\mathcal{F}$, F, that maximizes this quantity:
>
> $$\sum_{P_i \in P} \sum_{s \in P_i} p(s)^* \mathbf{A}(F(P_i), s)$$
>
> is:
>
> $$F(P_i) = \mathrm{Cond}(P_i) = p(\cdot \mid P_i)$$
>
> when **A** is strictly proper.

G&W can be used to derive Greaves and Wallace's claim about experiments:

> CONDMAX: Suppose you know that you are going to perform an experiment, $\mathcal{E}$. The update procedure that maximizes expected accuracy in response to $\mathcal{E}$, relative to probability function p, is the update-procedure that assigns, to each $\mathcal{E}_i$, $p(\cdot \mid \mathcal{E}_i)$.

The argument from G&W to CONDMAX**,** using our generalized framework, is simple.

> <u>Proof of CONDMAX:</u>
>
> (1) The expected accuracy of an update procedure U in response to an experiment $\mathcal{E}$, relative to a probability function p is:
>
> (*) $$\sum_{\mathcal{E}_i \in \mathcal{E}} \sum_{s \in \mathcal{E}_i} p(s)^* \mathbf{A}(U(\mathcal{E}_i), s)$$
>
> (from Lemma 2).

(2) The value of *U* that maximizes (*) is U=Cond($\mathcal{E}_i$).

(This follows from G&W and the fact that $\mathcal{E}$ is a partition)

(3) The update procedure *U* that maximizes expected accuracy in response to an experiment $\mathcal{E}$ is U=Cond($\mathcal{E}_i$). That is, the update procedure that maximizes expected accuracy is the one that has the agent conditionalize on the member of $\mathcal{E}$ that she learns.
(This follows from (1) and (2)).

But what about cases in which our future learning experiences aren't representable as experiments? Which update procedure maximizes expected accuracy in those cases? Here is the answer:

GENERALIZED CONDMAX**:** Suppose you know that you are going to undergo a learning experience, $\mathcal{X}$. The update procedure that maximizes expected accuracy in response to $\mathcal{X}$, relative to probability function *p*, is the update procedure that assigns, to each $\mathcal{X}_i$, $p(\cdot|\mathrm{L}(\mathcal{X}_i))$ where $\mathrm{L}(\mathcal{X}_i)$ is the proposition that $\mathcal{X}_i$ is the strongest proposition the agent exogenously learns upon undergoing the learning experience.

Proof of GENERALIZED CONDMAX:
Recall that the expected accuracy of an update procedure, *U*, in response to a learning experience $\mathcal{X}$ is defined as:

$$(\#) \qquad \sum_{\mathrm{L}(\mathcal{X}_i)\in \mathrm{L}(\mathcal{X})} \sum_{s\in \mathrm{L}(\mathcal{X}_i)} p(s)^* \mathbf{A}(U(\mathcal{X}_i)), s)$$

We are aiming to show that (#) is maximized when $U(\mathcal{X}_i) = \mathrm{Cond}(\mathrm{L}(\mathcal{X}_i))$. So suppose for reductio that this is false: that is, that there exists a function, U^*, such that:

$$\sum_{\mathrm{L}(\mathcal{X}_i)\in \mathrm{L}(\mathcal{X})} \sum_{s\in \mathrm{L}(\mathcal{X}_i)} p(s)^* \mathbf{A}(U^*(\mathcal{X}_i)), s) \quad > \quad \sum_{\mathrm{L}(\mathcal{X}_i)\in \mathrm{L}(\mathcal{X})} \sum_{s\in \mathrm{L}(\mathcal{X}_i)} p(s)^* \mathbf{A}(\mathrm{Cond}(\mathrm{L}(\mathcal{X}_i)), s)$$

Now, define $\mu(L(\mathcal{X}_i))$ as $U^*(\mathcal{X}_i)$.[13] It follows that:

$$\sum_{L(\mathcal{X}_i)\in L(\mathcal{X})} \sum_{s\in L(\mathcal{X}_i)} p(s)^* \mathbf{A}(\mu(L(\mathcal{X}_i)), s) \quad > \sum_{L(\mathcal{X}_i)\in L(\mathcal{X})} \sum_{s\in L(\mathcal{X}_i)} p(s)^* \mathbf{A}(\mathrm{Cond}(L(\mathcal{X}_i)), s)$$

But this is impossible, because it follows from G&W that the quantity:

$$(\#\#) \qquad \sum_{L(\mathcal{X}_i)\in L(\mathcal{X})} \sum_{s\in L(\mathcal{X}_i)} p(s)^* \mathbf{A}(F(L(\mathcal{X}_i)), s)$$

is maximized when $F(L(\mathcal{X}_i))= \mathrm{Cond}(L(\mathcal{X}_i))$. Thus, there cannot exist a μ that satisfies the inequality above. Contradiction.

Here is the lesson to be learned from CONDMAX and its generalization: the update procedure that maximizes expected accuracy in response to any learning experience is one in which an agent who learns $\mathcal{X}_i$ conditionalizes on the proposition *that she learns* $\mathcal{X}_i$ upon undergoing the learning experience.[14] The reason that conditionalizing on the proposition that one learns maximizes expected accuracy *in response to an experiment* is that, in these special cases, the agent knows that she will learn $\mathcal{X}_i$ if and only if $\mathcal{X}_i$ is true. In these cases, conditionalizing on $\mathcal{X}_i$ amounts to the very same thing as conditionalizing on $L(\mathcal{X}_i)$.[15]

---

[13] How do we know that there is such a μ? Since there is a bijection between the $\mathcal{X}_i$ and the $L(\mathcal{X}_i)$, there exists an inverse of $L(\mathcal{X}_i)$, which we'll call '$L^-(\mathcal{X}_i)$', such that $L^-(L(\mathcal{X}_i)) = \mathcal{X}_i$. We can then let $\mu(L(\mathcal{X}_i))$ be U^* composed with L^-. Thus: $\mu(L(\mathcal{X}_i)) = U^*(L^-(L(\mathcal{X}_i)) = U^*(\mathcal{X}_i)$.

[14] Note that this is true for any proposition that is the strongest proposition one exogenously learns, including propositions that are, themselves, about gaining information. So if, say, in a Monty Hall case, one thinks that the strongest proposition learned is something along the lines of: 'I gained the information that there is a goat behind door 2', the update procedure that maximizes expected accuracy will have you conditionalize on: '*I learned that* I gained the information that there is a goat behind door 2'.

[15] The result can be generalized further to cases in which the possible number of propositions learned is infinite. However, to perform this generalization, we need a notion of expected accuracy that doesn't rely on summation. Easwaran (2013) provides such a notion and argues, using this notion, that conditionalization maximizes expected accuracy. Like Greaves and Wallace, however, Easwaran relies on both PARTITIONALITY and FACTIVITY. So some modifications need to be made to derive GENERALIZED CONDMAX using Easwaran's framework. Since Easwaran's notion of expected accuracy is quite complex, I cannot, in this note, explain in general terms how the proof must be modified. But for those readers familiar with Easwaran's argument, here are the relevant details: First, Easwaran's claim that '*V* and *V'* are identical on ~E' (p.136) relies on FACTIVITY. For suppose FACTIVITY is violated. Then it's possible that, for some s, the agent learns E in s but ~E is true in s. In such a state $V(s) = I(A, x, s)$ and $V'(s) = I(A, x', s)$. Since it has not been assumed that x and x' are identical, it cannot be assumed that *V* and *V'* are identical on ~E. What can be assumed, however,

4. Iteration Principles

The update procedure that maximizes expected accuracy in general is not conditionalization. It is conditionalization*: conditionalizing on the proposition that one learned P, when P is the proposition learned.

Recall that we are interested in the expected accuracy of update procedures like conditionalization or conditionalization* because of the possibility that expected accuracy considerations can be used to support claims about which update procedures are rational. And recall that underlying the arguments under discussion for the rationality of various update procedures is the following assumption:

> RATACC*:* The rational update procedures are those that maximize expected accuracy according to a strictly proper scoring rule.

Together, RATACC and GENERALIZED CONDMAX entail:

> COND*: The rational update procedure is conditionalization*. In other words, upon learning P, an ideally rational agent will conditionalize on the proposition that she learned P.[16]

Since conditionalizing on any proposition involves assigning credence 1 to that proposition, and conditionalization* has us conditionalize on the proposition that we learned P, when P is learned, it follows from COND* that:

> LL: If one learns P, one is rationally required to be certain that one learned P.

I suspect that people who deny KK – the principle that whenever one knows P one is in a position to know that one knows P[17] – or related iteration principles, will find

---

without relying on FACTIVITY, is that *V* and *V'* are identical *on ~L(E)*. Second, Easwaran's claim that 'on E, *V*(s) = *I*(A, x, s) and *V'*(s) = *I*(A, x', s)' (p.136) relies on PARTITIONALITY. For suppose that PARTITIONALITY is violated. Then it's possible that there is some state s in which E is true but the agent doesn't learn E – rather, she learns some other proposition E*. In such a case, *V*(s) = *I*(A, f(E*), s) and *V'*(s) = *I*(A, f'(E*), s). Since it is not assumed that f(E*) is x, or that f'(E*) is x', we cannot assume that, on E, *V*(s) = *I*(A, x, s) and *V'*(s) = *I*(A, x', s). What can be assumed, however, without relying on PARTITIONALITY, is that, *on L(E)*, *V*(s) = *I*(A, x, s) and *V'*(s) = *I*(A, x', s). Plugging in these substitutions throughout the remainder of the proof yields the result that, in general, conditionalizing on L(E) (rather than E), where E is the proposition learned, is the update procedures that maximizes expected accuracy.

[16] Recall that the proposition one 'learns' refers to the strongest proposition one exogenously learns.

LL unattractive.[18] But if LL is rejected, COND* must also be rejected. In this section, I explore a number of ways of resisting the conclusion that conditionalization* is the rational update procedure, and the resulting commitment to LL. The most straightforward way to do this is to simply reject RATACC – the claim that the rational update procedures are those that maximize expected accuracy. Ultimately, I think that this is the most promising route for those who wish to reject COND* and/or LL. But first I'd like to describe two alternatives. The first involves claiming that all rational agents do, in fact, satisfy PARTITIONALITY and FACTIVITY. The second involves a modification of RATACC.

4.1 Endorsing the Requirements of PARTITIONALITY and FACTIVITY

The argument against the claim that conditionalization maximizes expected accuracy in general relied on the thought that rational agents may fail to satisfy PARTITIONALITY or FACTIVITY. I offered considerations that tell against the requirement that rational agents satisfy both of these conditions. But perhaps, upon realizing that endorsing conditionalization* as the rational update procedure brings with it a commitment to LL, one may want to revisit this issue.

However, even if a case can be made that all rational agents satisfy PARTITIONALITY and FACTIVITY, this won't help the LL-denier. For CONDMAX tells us that if all rational agents satisfy PARTITIONALITY and FACTIVITY, ordinary conditionalization will be the update procedure that maximizes expected accuracy. However, by Lemma 1, all rational agents who satisfy PARTITIONALITY and FACTIVITY will regard L(P) and P as equivalent. So, if rational agents conditionalize on P, upon learning P, they will assign credence 1 to P. But, since these agents assign credence 1 to P ↔ L(P), conditionalizing on P will result in the agent assigning credence 1 to L(P) as well. Thus, if PARTITIONALITY and FACTIVITY *are* satisfied, *conditionalization* yields the result that an agent that learns P will be certain that she learned P.[19]

---

[17] See, for example, Williamson (2000).
[18] Note, however, that at least some objections to KK don't extend to LL. KK has the consequence that if an agent knows P, she knows that she knows P, she knows that she knows that she knows P, and so on. However, recall that by 'learn' we mean exogenously learn. Thus, LL just says that if an agent *exogenously* learns P she must become certain that she exogenously learned P. It doesn't say that if she exogenously learns P, she *exogenously learns* that she exogenously learns P. The certainty in learning P need not, itself, be the result of exogenous learning. Thus, unlike KK, LL 'iterates' only once.
[19] Bronfman (2014) gives a related argument for the claim that agents that satisfy these conditions will conform to KK.

This brings out an important point: conditionalization and conditionalization* only yield different results when an agent doesn't satisfy at least one of PARTITIONALITY or FACTIVITY. I suggested that, in many ordinary cases, these requirements are not both satisfied. In such cases, conditionalization*, and not conditionalization, maximizes expected accuracy. But even if one disagrees with me about whether rational agents always satisfy PARTITIONALITY and FACTIVITY, one shouldn't reject the claim that conditionalization* maximizes expected accuracy. For conditionalization and conditionalization* amount to the very same thing when PARTITIONALITY and FACTIVITY are satisfied. Thus, COND* and LL follow from RATACC even if agents are rationally required to satisfy PARTITIONALITY and FACTIVITY.

4.2 Modifying RATACC

Aaron Bronfman (2014, p. 887-8) considers and rejects a rule that is similar to conditionalization*. His reason for rejecting the rule is based on the thought that when we're considering which update procedures maximize expected accuracy, we should only consider those procedures that the agent in question can competently execute. On this view, the rational update procedure isn't the update procedure *from the pool of possible update procedures* that maximizes expected accuracy. Rather, the rational update procedure is the procedure *from the pool of update procedures that the agent can competently execute* that maximizes expected accuracy.

As an example, suppose that Al fails to satisfy one of PARTITIONALITY or FACTIVITY. I have shown that the update procedure that maximizes expected accuracy for Al *from the pool of possible update procedures* is conditionalization*. But now suppose that Al sometimes exogenously learns P, but is unable to realize that he learned P. Arguably, Al can't competently execute conditionalization*.[20] If this is right, then according to modified RATACC, which has us consider only update procedures that Al can competently execute, Al is *not* required to conditionalize*.

I think that this is an interesting suggestion, but it is worth noting a few things: First, when we calculate the expected accuracy of update procedures, we always do so from the perspective of the agent prior to undergoing the learning experience. Bronfman's suggestion is that we remove from the pool of candidate update procedures those that the agent cannot execute. But what if the agent has

[20] How plausible this claim is depends on the modal scope of 'can'. I will simply assume that someone who is sympathetic to this line of thought will have a way of making sense of the modal that yields the desired result.

false (but rational) beliefs about which update procedures she can execute? Then the update procedure that maximizes expected accuracy from the pool of procedures that she *thinks* she can execute may differ from the update procedure that maximizes expected accuracy from the pool of procedures that she *can* execute. But it seems against the spirit of Bronfman's proposal to demand that the agent update in accord with the update procedure that maximizes expected accuracy relative to her *actual* abilities when she has no way of knowing which update procedure this is.

One might modify Bronfman's proposal so that what's relevant is not the agent's actual abilities, but the agent's opinions concerning her abilities. But if the only update procedures in the pool that she should be choosing from are those that she is *certain* that she will be able to execute, the pool may well be empty. Perhaps, then, the pool shouldn't only contain procedures that she is certain she will be able to execute. Maybe it should contain those procedures that she *believes* she can execute, or those that she is *sufficiently confident* that she can execute.[21] But there are additional complications. For suppose I now rationally believe that I won't be able to refrain from being certain that my child is the best player on the team, whatever evidence I receive. But I am wrong about this. In fact, I will perfectly well be able to evaluate the evidence concerning the relative abilities of my child. The view under consideration entails that even if, when the time comes, all of my evidence suggests that my child is mediocre, and I am capable of recognizing this fact, in virtue of the fact that, at an earlier time, I *believed* that I couldn't help but be certain that she is best, I am rationally *required* to be certain that she is the best! This seems highly implausible.

I don't mean to claim that these complexities are insurmountable, but it is worth noting that nothing that looks like ordinary conditionalization will emerge as a result of Bronfman's modification. If we modify RatAcc in the way Bronfman suggests and thereby avoid a commitment to Cond* and LL, what we are left with isn't good old-fashioned conditionalization. Rather, the rational update procedure will be something very messy and agent-relative that can't be neatly characterized in a formal framework. If we want to account for the limitations of non-ideal agents

---

[21] If we included only those procedures that the agent *knows* she can execute, then, since 'knows' is factive, we will run into the earlier problem. If the agent rationally believes that she can execute all of the procedures in set S, but she only *knows* that she can execute the procedures in S', then the view would imply that it's rational for her to accord with the update procedure that maximizes expected accuracy relative to S'.

this is to be expected, but we are now quite far from the project as Greaves and Wallace, and others involved in accuracy-first epistemology, originally conceived of it. In describing the idealized agents under discussion Greaves and Wallace say: 'Real epistemic agents are not (at least not quite) like this. Bayesian epistemology is a normative theory rather than a purely descriptive one' (p. 608). Greaves and Wallace are interested in a notion of *ideal* rationality that doesn't take an agent's cognitive limitations into account. One might have qualms about such idealizations, but these qualms will extend to Bayesian epistemology more generally and are not ones that I will address here.

Still, one might claim, even the idealized notion of rationality that Greaves and Wallace are working with takes into account some of the agent's limitations. After all, if *any* update procedure were allowed in the pool, then surely the update procedure that maximizes expected accuracy would be one that requires that, in every state, the agent assign credence 1 to all the truths and credence 0 to all the falsehoods!

Now, as a matter of fact, given the way we have defined 'update procedure', the rule 'assign credence 1 to all truths and 0 to all falsehoods' (let's call it 'the truth rule') simply isn't an update procedure. For recall that an update procedure is just a function from the propositions one might learn to credence functions. Since the truth values of some propositions may vary amongst the worlds in which the agent learns the same information, but the recommended credence function cannot vary amongst these worlds, a function from the propositions the agent might learn to credence functions will not, in general, be one that, when conformed to, results in an agent assigning credence 1 to all truths in every state.

Nonetheless, one might think that the reason we defined the notion of an update procedure in a way that rules out the truth rule is that we are only interested in procedures that are, in some sense, *available* to the agent upon undergoing the learning experience. We don't want to require that the agent be certain that it will rain tomorrow, in virtue of the fact that it *will* rain tomorrow, if all she learns is, say, that a coin landed Heads. Similarly, you might think, we must find *some* way of ruling out update procedures that require an agent to be certain that she learned P, in virtue of that fact that she *did* learn P, even though the only information she exogenously received was that P. Perhaps so. But the issue here will be: available *in what sense*?

It will be helpful to make use of Ned Hall's (2004) distinction between analyst experts and database experts. We defer to database experts because they possess a great deal of information. We defer to analyst experts because of their superior information processing abilities. Thus, we can distinguish agents who are idealized along the database dimension (they are certain of all and only the truths), and agents who are idealized along the analyst dimension. It is the latter kind of idealization that Greaves and Wallace are interested in. They want to know how an idealized analyst will revise her beliefs in light of new information. Since they are interested in idealized information *processing*, and not idealized information *possession*, it is clear why they require that update procedures issue the same recommendations in any two states in which the agent gains the same information. It will, however, be difficult to come up with a principled way of ruling out conditionalization* as the ideal update procedure if the ideal in question is ideal information processing. This is because, like conditionalization, conditionalization* is simply an operation performed on the proposition exogenously learned. The operation is the following: If P is the proposition learned, take P, attach an L to it, and conditionalize on the resulting proposition: L(P).

If we were happy with ordinary conditionalization, then we were happy with requiring that (ideal!) agents be certain that Q, upon learning P, if P entails Q. In endorsing this commitment we needn't suppose that any event in which one exogenously learns P *constitutively* involves a learning of Q. Rather, the Q-learning may be a kind of endogenous learning that idealized agents will undergo upon exogenously learning P. The requirement that agents be certain that they learned what they learned is, in the relevant sense, no different from the requirement that agents be certain in the propositions that their evidence entails. Here too, we needn't suppose that any event in which one learns P *constitutively* involves a learning that one learned P. The claim is rather that *ideal* agents will come to be certain that they learned P upon appropriate processing of the information that P.

In sum, Bronfman's modification of RatAcc may well be worth serious consideration, but it does not engage with the project as Greaves and Wallace conceived of it: figuring out the ideally rational way to revise beliefs, where the idealization in question is along the information processing dimension. Conditionalization*, just like conditionalization, is an operation on the proposition an agent learns. If we are interested in ideal information processing, there shouldn't be any restrictions on what operations can be performed on this proposition.

4.3 Giving up RatAcc

Rather than trying to modify RATACC, one may simply reject the idea that anything in the vicinity of RATACC is true. On this view, there simply is no straightforward connection between the rational way of revising one's credences and considerations of expected accuracy.

There is plenty of literature devoted to evaluating the merits of accuracy-first epistemology[22] and entering into this debate will take us beyond the scope of this paper. But it is important to realize that RATACC plays an important role in much of the accuracy-first project.[23] So I will simply note that this is one way that someone who wants to reject COND* and LL might go. If this turns out to be the only acceptable way to reject these claims, then we will have learned the following interesting fact: Many accuracy-first epistemologists (those who endorse RATACC) are committed to some substantive iteration principles, and, conversely, those who reject such principles are committed to rejecting large portions of the accuracy-first project.

5. Further Generalizations and Further Consequences

I have shown that conditionalizing on the propositions we learn does not, in general, maximize expected accuracy. Rather, conditionalizing on the proposition that we learned P, when P is the proposition learned, is the update procedure that maximizes expected accuracy. In this section I provide further generalizations of this result and show that, no matter which features of an agent's situation the RATACC-er thinks the rationality of an agent's credence function depends on, she is committed to:

> LUMINOUS INFALLIBILITY: There is a class of propositions concerning an agent's situation, such that, for any subject S, if S is rational, these propositions will be true of S if and only if she is certain of them.

To begin, I will give an argument for the following generalization of GENERALIZED CONDMAX:

[22] See, for example, Caie (2013), Greaves (2013), Pettigrew (2016), Konek and Levinstein (forthcoming), and Carr (ms.).

[23] Though see Schoenfield (forthcoming), section 4, for an alternative conception of how rationality and accuracy considerations interrelate, which takes accuracy as fundamental, but gives up on RATACC.

SUPER GENERALIZED CONDMAX: Let U be a function from a set of propositions $\mathcal{X}$ to credence functions with the intended interpretation that an agent conforming to U adopts $U(\mathcal{X}_i)$ whenever the agent bears relation R to $\mathcal{X}_i$. Let R($\mathcal{X}_i$) be the proposition 'The agent bears relation R to $\mathcal{X}_i$'. If the R($\mathcal{X}_i$) form a partition, then the function, U, such that conforming to U maximizes expected accuracy, is the one that has the agent conditionalize on the proposition 'the agent bears relation R to $\mathcal{X}_i$' whenever the agent bears relation R to $\mathcal{X}_i$.

Why is this principle true? Here's the intuitive idea: Suppose you could choose a credence function that you knew an agent would adopt whenever she bears the relation R to some proposition P. Even without knowing anything else about this relation, if you wanted her to be as accurate as possible, the following seems like a sensible first step: have her assign credence 1 to the proposition: *she bears relation R to P* whenever she bears relation R to P. For this will guarantee that if she conforms to the procedure, she will assign credence 1 to a truth! What conditionalizing on 'she bears relation R to P' adds to this is just that she'll renormalize the rest of her credences in response to her newfound certainty.

More formally, note that SUPER GENERALIZED CONDMAX differs from GENERALIZED CONDMAX only in that we are talking about bearing the R relation to a proposition, rather than the learning relation, and we are understanding what it is for an agent to conform to U as the agent adopting $U(\mathcal{X}_i)$ whenever she bears R to $\mathcal{X}_i$ instead of whenever she learns $\mathcal{X}_i$. But the proof of GENERALIZED CONDMAX didn't rely on any special feature of L. So if, instead of asking how an agent should revise her credences as a function of which proposition she *learns*, we ask how an agent should revise her credences as a function of which proposition she *bears R to*, it will turn out that the update procedure that maximizes expected accuracy is the one that has the agent conditionalize on R($\mathcal{X}_i$) whenever she bears R to $\mathcal{X}_i$.

SUPER GENERALIZED CONDMAX explains why the results in this paper are completely neutral with respect to one's understanding of 'learning'. A theorist can take any notion of learning that she's interested in (coming to assign credence 1 to P, coming to know P, coming to believe P), and partition an agent's possibility space in accord with the different 'learnings' she might undergo (perhaps including a trivial instance of learning to capture a case in which no new information is gained). Then, SUPER GENERALIZED CONDMAX will say that the update procedure that maximizes expected accuracy in response to whatever kind of learning takes place is conditionalizing on the proposition *that the relevant kind of learning has taken place.*

We can now generalize the result even further. For it also follows from G&W that:

SUPER-DUPER (SD) GENERALIZED CONDMAX**:** Consider any partition of propositions P_i over a set of states Ω. Let *U* be a function from P_i to credence functions with the intended interpretation that an agent adopts $U(P_i)$ whenever P_i obtains. The *U* that maximizes expected accuracy is the one that assigns to each P_i the credence function that results from conditionalizing on P_i.

<u>Proof</u>
Let P_i be a set of propositions that partition Ω. It follows directly from G&W that the function *U* that maximizes this quantity:

$$\sum_{P_i \in \Omega} \sum_{s \in P_i} p(s) * \mathbf{A}(U(P_i), s))$$

is: U=Cond(P_i).

Note that the quantity above represents the expected accuracy of an agent's credences if that agent adopts $U(P_i)$, whenever P_i obtains. It follows that if we understand conforming to *U* as adopting $U(P_i)$ whenever P_i obtains, the *U* which is such that conforming to it maximizes expected accuracy is the one that has the agent adopt Cond(P_i) whenever P_i obtains.

This result allows us even greater flexibility in terms of how we think about exogenously gaining information because we are no longer required to think of information gaining as an agent coming to bear a relation to a *proposition*. Suppose, for example, that one thought that the world doesn't fling single propositions at agents, but *sets* of propositions. We then might ask: how should one revise one's credences in response to learning a set of propositions? SD-GENERALIZED CONDMAX entails that the update procedure that maximizes expected accuracy in response to learning sets of propositions is the one that has the agent conditionalize on the proposition 'S is the set of propositions that was learned' whenever S is the set of propositions learned. Or suppose that, like Richard Jeffrey (1992), one thinks that exogenous learning involves the world shifting around some of an agent's credences. We then might ask: how should one revise one's credences in response to certain *credal changes* taking place? If we're looking for the update procedure in response to credal changes that maximizes expected accuracy, the answer will *not* be to Jeffrey-conditionalize. The answer will be to regular-old-conditionalize on the proposition 'such and such credal changes have occurred' whenever such and such credal changes have occurred.

There is a sense, then, in which a defender of RATACC can't help but adopt *some version* of the truth rule. For whatever one's theory of rationality is, one can partition the space of possible situations an agent might find herself in in such a way that the same doxastic state is rational in each cell of the partition. Perhaps, for example, a theorist partitions the space based on what the agent's phenomenology is: {She has phenomenology P1, she has phenomenology P2…} or what she learns {She learns $\mathcal{X}_1$, she learns $\mathcal{X}_2$…} or what her evidence is: {She possesses E1, She possess E2…}. Call this partition, whatever it is, P. There is, then, a function from the $P_i \in P$ to credence functions that (for this theorist) represents the credence function that is rational for an agent to adopt in any given cell of the partition. We'll call the propositions in P *the propositions whose truth determines what credence function it is rational for an agent to adopt*.

Now, suppose that our theorist is a RATACC-er. It follows from SD-GENERALIZED CONDMAX that conditionalizing on P_i whenever P_i obtains is the way to assign credence functions to the members of P that maximizes expected accuracy. So the RATACC-er will think that if P_i is true, a rational agent will conditionalize on P_i and so become certain that it is true. The RATACC-er must also think that if a rational agent is certain that P_i, then P_i is true. This is because the P_i form a partition, and so a rational agent will be certain of at most one P_i. (If she were certain of more than one P_i, then she would be certain of two incompatible propositions). We also know that she will be certain of at least one P_i, since at least one P_i will be true, and we already established that if P_i is true she will be certain that it is (since she will have conditionalized on it). It follows that she will be certain of *exactly* one P_i: the true one. Thus, for any P_i, the agent will be certain that P_i, if and only if P_i is true. In other words:

> If RATACC is true, then the propositions whose truth determines what credence function it is rational for an agent to adopt are propositions that a rational agent is luminously infallible about – that is, they are propositions that she will be certain of if and only if they are true.

We are now in a better position to recognize the awkwardness that arises in the Greaves and Wallace framework – an awkwardness that, I believe, reflects a tension in our thinking about these issues more generally. In defining an update procedure, Greaves and Wallace committed themselves to the view that which credence function it is rational for an agent to adopt depends on which proposition the agent learns. In other words, for Greaves and Wallace, the P_i – those propositions whose *truth* determines which credence function it is rational for an agent to adopt – are the $L(\mathcal{X}_i)$: the propositions describing which proposition an agent learns. While this seems like a perfectly plausible choice for one's P_i, it would *not* be plausible to suppose that the credence function that it is rational for an agent to adopt depends on whether, for example, it's raining in Singapore, regardless of what evidence the agent has to that effect. Thus, propositions about the weather conditions in Singapore are *not* a plausible choice of P_i. But here's the problem: Intuitively, we think that propositions about the weather in Singapore are propositions we might learn, and we all grew up liking the idea that conditionalizing on *what we learn* is rational. But what follows from SD-GENERALIZED CONDMAX and RATACC is that whichever propositions are such that their *truth* determines which credence function is rational are the propositions that a rational agent will conditionalize on. So a RATACC-er can't think that which of the $L(\mathcal{X}_i)$ is true, determines which credence function is rational, *and* think that the propositions that the agent will conditionalize on are propositions about the weather, unless she also thinks that the propositions about the weather are equivalent to the propositions in $L(\mathcal{X})$. This is why, to get the Greaves and Wallace result – that conditionalizing on *the content* of what one learns maximizes expected accuracy – we must impose such severe restrictions on the possible contents of learning. For whatever these contents are, they must be equivalent from our perspective to the proposition that we learn them.

6. Conclusion

I have argued that conditionalization is not the update procedure that, in general, maximizes expected accuracy. The update procedure that maximizes expected accuracy is conditionalization*: conditionalizing on the proposition that one learned P when P is the strongest proposition one exogenously learned. Conditionalizing on P, it turns out, only maximizes expected accuracy in cases in which the agent is antecedently certain that, for all P she might learn, if P is true she will learn it, and if she learns P, it is true.

If the rational update procedures are those that maximize expected accuracy (that is, if RatAcc is true), the fact that conditionalization* maximizes expected accuracy entails that conditionalization* is rational, and if conditionalization* is rational, then one is rationally required to be certain that one learned P whenever it is true that one learned P.

These results are instances of a yet deeper phenomenon. Anyone who accepts RatAcc is committed to the existence of a class of propositions that rational agents will be luminously infallible about: a class of propositions that rational agents will be certain of if and only if they are true.

The results in this paper can thus be summarized as follows: It follows from RatAcc that:

(1) If which credence function it is rational to adopt is determined by which proposition one learns, then conditionalizing on the proposition *that one learned* $\mathcal{X}_i$, when $\mathcal{X}_i$ is the proposition learned, is the rational way of revising one's credences. The class of propositions $L(\mathcal{X}_i)$ will be the class of propositions that a rational agent is luminously infallible about.

(2) If a rational agent regards any proposition $\mathcal{X}_i$ that she might learn as equivalent to $L(\mathcal{X}_i)$, the claim that she learned it, then the rational update procedure (conditionalizing on $L(\mathcal{X}_i)$) will amount to the same thing as conditionalizing on $\mathcal{X}_i$. In this case, the class of propositions *one might learn* (the $\mathcal{X}_i$) are also propositions that a rational agent is luminously infallible about.

(3) If which credence function it is rational to adopt is determined by some other feature of an agent's situation, such that, for some partition P, which credence function it is rational for an agent to adopt depends on which member of P is true, then a rational agent will conditionalize on P_i (a member of P) whenever P_i is true. The propositions $P_i \in P$ will be the propositions that a rational agent is luminously infallible about.

Committed Rat-Accers might take these results as favoring a kind of foundationalist epistemology on which there is some privileged class of propositions that rational agents will be certain of if and only if they are true. Adamant deniers of such an epistemology might take these arguments as a reason to abandon the idea that the rational update procedures are those that maximize expected accuracy. But however we proceed, it is important to be aware of the extent to which the thought that rationality involves maximizing expected accuracy and such claims as LUMINOUS INFALLIBILITY are intertwined. They will, I believe, stand or fall together.[24]

References

Arnold, A. (2013). 'Some Evidence is False', *Australasian Journal of Philosophy* 91 (1): 165-72

Bronfman, A. (2014). 'Conditionalization and not Knowing that one Knows', *Erkenntnis* 79(4): 871-92

Caie, M. (2013). 'Rational Probabilistic Incoherence', *Philosophical Review* 122(4): 527-75

Carr, J. (ms.). 'Epistemic Utility Theory and the Aim of Belief'

Comesaña, J. and McGrath, M. (2014). 'Having False Reasons' in Littlejohn, C. and Turri, J. (eds.) *Epistemic Norms*. Oxford University Press

Drake, J. (forthcoming). 'Motivating Reason to Slow the Factive Turn in Epistemology' in Mitova, V. (ed.). *The Factive Turn in Epistemology*. Cambridge University Press

Easwaran, K. (2013). 'Expected Accuracy Supports Conditionalization – and Conglomerability and Reflection', *Philosophy of Science* 80 (1): 119-142

Greaves, H. and Wallace, D. (2006). 'Justifying conditionalisation: conditionalisation maximizes expected epistemic utility', *Mind* 115(459): 607-32

Greaves, H. (2013). 'Epistemic Decision Theory', *Mind* 122(488): 915-52

Hall, N. (2004). 'Two Mistakes about Credence and Chance', *Australasian Journal of Philosophy*. 82(1): 93-111

Horowitz, S. (2013). 'Immoderately Rational', *Philosophical Studies* 167(1):1-16

Howson, C. and Urbach, P. (1989). *Scientific Reasoning: The Bayesian Approach*. Open

[24] I am grateful to audiences at the Massachusetts Institute of Technology, Rutgers University, the University of Bristol's epistemic utility theory project, and The University of Texas at Austin for extremely helpful feedback on this paper. Special thanks to Branden Fitelson, Joshua Dever, Sinan Dogramaci, Brian Knab, Maria Lasonen Aarnio, David Sosa and especially Susanna Rinard for helpful discussion and for reading and commenting on earlier drafts. Thanks also to the anonymous referees at *Mind* for helpful comments and suggestions. Some of the work for this paper was conducted at Oxford University and supported by a grant to the *New Insights and Directions for Religious Epistemology* project, funded the John Templeton foundation.

Court

Hutchison, K. (1999). 'What are Conditional Probabilities Conditional Upon?', *British Journal for the Philosophy of Science* 50: 665-95

Jeffrey, R. (1992). *Probability and the Art of Judgment*. Cambridge University Press

Joyce, J. (2009). "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief" in F. Huber & C. Schmidt-Petri (eds.). *Degrees of Belief*. Synthese Library. 263-297

Konek, J. and Levinstein, B. (forthcoming). 'The Foundations of Epistemic Decision Theory', *Mind*

Miller, B. (forthcoming). 'How to Be a Bayesian Dogmatist', *Australasian Journal of Philosophy*

Moss, S. (2011). 'Scoring Rules and Epistemic Compromise', *Mind* 120(480):1053 -69

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

Rizzierie, A. (2011). 'Evidence does Not Equal Knowledge', *Philosophical Studies* 153(2): 235-242

Schoenfield, M. (forthcoming). 'Bridging Rationality and Accuracy', *The Journal of Philosophy*

Teller, P. (1976). 'Conditionalization, observation, and change of preference' in Harper, W. and Hooker, C.A. (eds.). *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science.* D. Reidel

van Fraassen, B. (1989). *Laws and Symmetry*. Oxford University Press

van Fraassen, B. (1999). 'A New Argument for Conditionalization', *Topoi* 18: 93-96

Williams, P. (1980). 'Bayesian Conditionalization and the Principle of Minimum Information', *British Journal for the Philosophy of Science* 31(2): 131-44

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