mschoenfield@nyu.edu

Conditionalization Does Not (in general) Maximize Expected Accuracy

1. Background

Rational agents revise their opinions in light of new information they receive. We can think of information-processing as occurring in two stages: exogenous and endogenous.

The Question: How should agents revise their beliefs in light of the information they gain *exogenously*?

Bayesian Answer: By conditionalization. You conditionalize on E if

$$p_{new}(\cdot) = p_{\text{old}}(\cdot \mid \text{E})$$
where
$$p(\text{A}|\text{B}) = p(\text{A\&B}) \,/\, p(\text{B}).$$

Since conditionalizing is an operation performed on a *proposition*, thinking of conditionalizing as a way of responding to new information requires thinking that each body of information can be uniquely characterized by a proposition, and that in gaining information one comes to bear some relation to that proposition. I will use the term "exogenously learn" or "learn" for short to describe whatever this relation is.

Why conditionalize? Greaves and Wallace argue for the claim that *conditionalization maximizes expected accuracy*. Their argument for the rationality of conditionalization relies on:

RatAcc: The rational update-procedures are those that maximize expected accuracy.

> *Thesis 1*: If RatAcc is true, then the rational update-procedure in general is conditionalization*, and not conditionalization.

Conditionalization* has us conditionalize on the proposition *that we learn P*, when P is the proposition we learn.

Thesis 2 (Luminous Infallibility): If RatAcc is true then there exists a (nontrivial) set of propositions that a rational agent is luminously infallible about – that is, a set of propositions that a rational will be certain of if and only if they are true.

2. Formal Framework

Accuracy is measured by a scoring rule, **A,** which takes a credence function, *c*, and a state of the world, 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.

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. So this is a partition of your possibility space:

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

As are:

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.

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.

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**. An agent who knows she'll undergo a learning-experience can represent her future learning-experience by the set of propositions that she assigns non-zero credence to exogenously learning.

An **update-procedure,** *U*, in response to a future learning-experience, $\mathcal{X}$, is a function that assigns to each member of $\mathcal{X}$, a credence function, with the intended interpretation that an agent performing this update-procedure adopts $U(\mathcal{X}_i)$ as her credence function if and only if she learns $\mathcal{X}_i$.

Let $\mathbf{A}(U(s),s)$ represent the *accuracy score of an agent conforming to U in s.*

The **expected accuracy of an update-procedure *U*** in response to a learning-experience $\mathcal{X}$, relative to a probability function *p* is:

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

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

This quantity represents, roughly, how accurate we expect to be as a result of conforming to the update procedure.

3. Greaves and Wallace's Assumptions

> 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.

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

Is plausible that all rational agents satisfy both of these conditions?
Not obviously: You might think that I could rationally find myself in a position in which I'm certain that I'll learn exactly one of: {P, Q, P&Q}, and I assign non-zero credence to each. (In this case PARTITIONALITY fails).

You also might think that I could find myself in the position in which I leave open the possibility of becoming misinformed – that is, I leave open the possibility that the world will "fling" a false proposition into my belief box. (In this case FACTIVITY fails)

4. Three Theorems

G&W: Take any partition of states $\mathcal{P}$: $\{\mathcal{P}_1...\mathcal{P}_n\}$ and consider the set of functions, $\mathcal{F}$, that assign members of $\mathcal{P}$ to probability functions. The member of $\mathcal{F}$, F, that maximizes this quantity:

$$\sum_{\mathcal{P}_i \in \mathcal{P}} \sum_{s \in \mathcal{P}_i} p(s) * \mathbf{A}(F(\mathcal{P}_i), s)$$

is: $F(\mathcal{P}_i)$ = Cond = $p(\cdot \mid \mathcal{P}_i)$

CondMax: The update-procedure that maximizes expected accuracy relative to a probability function p that satisfies PARTITIONALITY and FACTIVITY is the update-procedure that assigns to each $\mathcal{X}_i$ that the agent thinks she might learn: $p(\cdot \mid \mathcal{X}_i)$.

Generalized CondMax**:** The update-procedure that maximizes expected accuracy relative to any probability function p is the update-procedure that assigns to each $\mathcal{X}_i$ that the agent thinks she might learn: $p(\cdot|L(\mathcal{X}_i))$, where $L(\mathcal{X}_i)$ is the proposition that the agent learns $\mathcal{X}_i$.

5. Three Consequences

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.

Cond* entails:

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

Super 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.

Luminous Infallibility: 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.