

Bayesian updating when what you learn might be false (forthcoming in *Erkenntnis*)*

Richard Pettigrew

December 1, 2020

Abstract

Michael Rescorla (2020) has recently pointed out that the standard arguments for Bayesian Conditionalization assume that whenever I become certain of something, it is true. Most people would reject this assumption. In response, Rescorla offers an improved Dutch Book argument for Bayesian Conditionalization that does not make this assumption. My purpose in this paper is two-fold. First, I want to illuminate Rescorla's new argument by giving a very general Dutch Book argument that applies to many cases of updating beyond those covered by Conditionalization, and then showing how Rescorla's version follows as a special case of that. Second, I want to show how to generalise R. A. Briggs and Richard Pettigrew's Accuracy Dominance argument to avoid the assumption that Rescorla has identified (Briggs & Pettigrew, 2018). In both cases, these arguments proceed by first establishing a very general reflection principle.

Careful formulations of the Bayesian norm of Conditionalization acknowledge that it governs how you should plan to update your credences, or how you should be disposed to update them. It does not govern how you should in fact update, or at least not directly. That is, Conditionalization does not govern the relationship between your prior credences and your posterior credences, but rather the relationship between your prior credences and your plans or dispositions for updating those priors. In particular, it governs those plans or dispositions you have for updating your credences in response to a certain sort of learning situation, namely, one in which you become certain of a proposition.

An updating plan comprises two components:

*Acknowledgements: I am extremely grateful to Michael Rescorla, as well as the referees for this journal, for extremely helpful and incisive comments on two earlier versions. The philosophy and the mathematics both improved as a result of heeding their suggestions.

- (i) a partition of the space of possibilities;
- (ii) a function that assigns to each element of that partition the posterior credence function that this plan endorses as a response to that element.

When we consider Conditionalization—and we therefore think of someone with a prior credence function c at time t , who is planning how to respond to evidence they receive between t and a later time t' that makes them certain of something—the partition consists of the following propositions:^{1,2}

I become certain of E_1 and nothing stronger by t' .

I become certain E_2 and nothing stronger by t' .

⋮

I become certain of E_n and nothing stronger by t' .

And the updating plan consists of conditional intentions or dispositions of the form:

If I become certain of E_1 and nothing stronger by t' , then I'll adopt credence function c_1 at t' .

If I become certain of E_2 and nothing stronger by t' , then I'll adopt credence function c_2 at t' .

⋮

If I become certain of E_n and nothing stronger by t' , then I'll adopt credence function c_n at t' .

Throughout, I assume that E_1, \dots, E_n are mutually exclusive.³ Conditionalization then says this:

Conditionalization Rationality requires that, if

- (i) my credence function at t is c ,
- (ii) c is defined on the algebra \mathcal{F} ,
- (ii) my updating plan contains: *If I become certain of E_i and nothing stronger by t' , then I'll adopt credence function c_i at t' , and*

¹I assume throughout that the partition is finite.

²Here, and in the remainder of the paper, I will talk of updating plans to cover both plans that you explicitly and consciously formulate as well as plans that describe how you are disposed to update.

³Rescorla (2020) assumes that they are exhaustive as well as exclusive, but it turns out that isn't necessary.

(iii) $c(E_i) > 0$;

then, for all X in \mathcal{F} ,

$$c_i(X) = c(X|E_i) := \frac{c(XE_i)}{c(E_i)}$$

That is, at t , I should make a plan to do the following at t' : to condition my prior credence function c on the strongest proposition of which I become certain at t' , providing my prior assigned positive credence to that proposition.

In favour of this version of Conditionalization, there are at least four arguments: two pragmatic, two epistemic. The first pragmatic argument is David Lewis's diachronic Dutch Book argument (Lewis, 1999); the second is Peter M. Brown's expected utility argument (Brown, 1976); the first epistemic argument is Hilary Greaves and David Wallace's expected accuracy argument (Oddie, 1997; Greaves & Wallace, 2006); the second is R. A. Briggs and Richard Pettigrew's Accuracy Dominance argument (Briggs & Pettigrew, 2018). However, as Michael Rescorla (2020) has recently noted, each of these arguments relies on an assumption, namely, that the strongest proposition of which you become certain by t' is true. That is, you never become certain of a proposition that is false—if you become certain that it is raining, then it is raining. As a result, if this assumption doesn't always hold, then these arguments don't establish Conditionalization. They establish only a restricted version of it that applies in those cases where the assumption holds.

Rescorla offers an improved Dutch Book argument for Conditionalization that does not assume that, if you become certain of E_i , then E_i must be true. My purpose in this paper is two-fold. First, I want to illuminate Rescorla's new argument by giving a Dutch Book argument for a general version of the Reflection Principle that applies to many cases of updating beyond those covered by Conditionalization, and then showing how Rescorla's version follows as a special case of that. Second, I want to show how to generalise Briggs and Pettigrew's Accuracy Dominance argument to avoid the assumption that Rescorla has identified. Indeed, I offer an accuracy dominance argument for the same general version of the Reflection Principle for which I offer a Dutch Book argument.

1 Two general norms connecting priors and possible posteriors

Van Fraassen's generalization of his original Reflection Principle says that, for each proposition X to which you currently assign a credence, your current credence in X should lie in the span of your possible future credences

in X (van Fraassen, 1995, 1999). That is, your current credence in a particular proposition should be a weighted average or convex combination of your possible future credences in that combination. So, if your current credence in X is x and your possible future credences in X are x_1, \dots, x_n , then there should be weights $0 \leq \lambda_1, \dots, \lambda_n \leq 1$ such that $\sum_{i=1}^n \lambda_i = 1$ and $x = \sum_{i=1}^n \lambda_i x_i$. Jonathan Weisberg (2007) calls this General Credence Reflection.

I want to consider a slightly stronger principle: your current credence function should lie in the span of your possible future credence functions. That is, if your current credence function is c and your possible future credence functions are c_1, \dots, c_n , then there should be $0 \leq \lambda_1, \dots, \lambda_n \leq 1$ such that $\sum_{i=1}^n \lambda_i = 1$ and $c = \sum_{i=1}^n \lambda_i c_i$.⁴ This is a more general norm than Conditionalization. Conditionalization covers cases in which you arrive at your future credence function in a particular way. It applies when you become certain of something and you set your future credence function by using your updating plan to respond to that. The principle that interests me covers those cases and other cases besides: cases in which your evidence shifts your opinion without making you certain of anything new, such as in Richard Jeffrey's case of viewing cloth by candlelight or van Fraassen's Judy Benjamin case (Jeffrey, 1965; van Fraassen, 1981).⁵

There is a weaker and a stronger version of this principle. The one I sketched in the previous paragraph is the weaker one. Here it is in more detail. Suppose c is your credence function at t , and R is the set of credence functions you might have at time t' . That is, c is your prior and R contains all and only your possible posteriors. We'll assume throughout that R is finite. Then the Weak General Reflection Principle says that c should be in the convex hull of R .⁶ That is, c should be a convex combination or weighted sum or weighted average of the credence functions in R . That is:

Weak General Reflection Principle Suppose c is your credence

⁴To see that this principle is indeed stronger:

(i) if c is in the span of c_1, \dots, c_n , then $c = \sum_{i=1}^n \lambda_i c_i$, so $c(X) = \sum_{i=1}^n \lambda_i c_i(X)$, and therefore $c(X)$ is in the span of $c_1(X), \dots, c_n(X)$;

(ii) suppose

$$\begin{array}{rcl} c(X) & = & 0.3 \\ c_1(X) & = & 0 \\ c_2(X) & = & 1 \end{array} \quad \begin{array}{rcl} c(Y) & = & 0.6 \\ c_1(Y) & = & 0 \\ c_2(Y) & = & 1 \end{array}$$

then $c(X)$ is in the span of $c_1(X), c_2(X)$, and $c(Y)$ is in the span of $c_1(Y), c_2(Y)$, but c is not in the span of c_1, c_2 . After all, for any $0 \leq \lambda_1, \lambda_2 \leq 1$ with $\lambda_1 + \lambda_2 = 1$, $(\lambda_1 c_1 + \lambda_2 c_2)(X) = \lambda_2 = (\lambda_1 c_1 + \lambda_2 c_2)(Y)$, whereas $c(X) \neq c(Y)$.

⁵van Fraassen (1999), Skyrms (1997), and Huttegger (2013) have also argued that reflection principles are generalizations of Conditionalization.

⁶The convex hull of a set is the smallest convex set that contains it. A set is convex when it contains any mixture of any two elements it contains. Thus, the interior of a circle is convex, but its circumference is not.

function at t , and $R = \{c_1, \dots, c_n\}$ is the set of credence functions you might have at t' . Then rationality requires that there is, for each c_i in R , a non-negative weighting $0 \leq \lambda_i \leq 1$ such that $\sum_{i=1}^n \lambda_i = 1$ and

$$c(-) = \sum_{i=1}^n \lambda_i c_i(-)$$

And the Strong General Reflection Principle says that your prior should be in the interior of the convex hull of your possible posteriors. That is, it should be a convex combination or weighted sum or weighted average of your possible posteriors *where the weights are all positive*. That is:

Strong General Reflection Principle Suppose c is your credence function at t , and $R = \{c_1, \dots, c_n\}$ is the set of credence functions you might have at t' . Then rationality requires that there is, for each c_i in R , a positive weighting $0 < \lambda_i < 1$ such that $\sum_{i=1}^n \lambda_i = 1$ and

$$c(-) = \sum_{i=1}^n \lambda_i c_i(-)$$

2 From General Reflection to Conditionalization

As we will see, we can establish Weak GRP via Dutch Book or Accuracy Dominance arguments; and we can establish Strong GRP via a Dutch Book argument, but not via an Accuracy Dominance argument. In this section, I show that we can obtain Conditionalization from either of these. This is not a particularly novel point. It was noted by van Fraassen when he introduced his general version of the Reflection Principle (van Fraassen, 1999).

Throughout, we will assume that all credence functions are defined on the same finite algebra of propositions, \mathcal{F} .

Now, suppose first that, between t and t' , you will become certain of exactly one of E_1, \dots, E_n , where each E_i is in \mathcal{F} . And suppose you have an updating plan defined on each of these possibilities. That is, the partition on which your updating plan is defined consists of the following propositions:

I become certain of E_i and nothing stronger by t' ,

for $i = 1, \dots, n$. And the plan itself consists of the following conditionals:

If I become certain of E_i and nothing stronger by t' , then I'll adopt credence function c_i at t' ,

for $i = 1, \dots, n$. And suppose finally that your possible future credence functions are the result of following your updating plan. So $R = \{c_1, \dots, c_n\}$. Now, suppose that

- (i) E_1, \dots, E_n are mutually exclusive propositions, and
- (ii) $c_i(E_i) = 1$, for $i = 1, \dots, n$.

Then, for $j = 1, \dots, n$,

$$\begin{aligned}
 c(E_j) &= \sum_{i=1}^n \lambda_i c_i(E_j) && \text{by Weak GRP} \\
 &= \lambda_j c_j(E_j) && \text{since } E_1, \dots, E_n \text{ are exclusive,} \\
 &&& \text{so } c_i(E_j) = 0 \text{ if } i \neq j \\
 &= \lambda_j && \text{since } c_j(E_j) = 1
 \end{aligned}$$

Thus, for all X in \mathcal{F} and E_j

$$\begin{aligned}
 c(XE_j) &= \sum_{i=1}^n \lambda_i c_i(XE_j) && \text{by Weak GRP} \\
 &= \sum_{i=1}^n c(E_i) c_i(XE_j) && \text{by above} \\
 &= c(E_j) c_j(XE_j) && \text{since } c_i(XE_j) = 0 \text{ if } i \neq j \\
 &= c(E_j) c_j(X) && \text{since } c_j(E_j) = 1
 \end{aligned}$$

Thus, if $c(E_j) > 0$, then $c_j(X) = c(X|E_j)$. So the future credence function c_j you plan to adopt if you become certain of E_j is obtained from your prior c by conditioning it on E_j , providing c assigns positive probability to E_j . This is what we wished to show. And note: at no point did we assume that you take yourself to learn E_j only if E_j is true.

3 The Dutch Book argument for GRP

In the Dutch Book argument for Probabilism, we show that, if your credence function does not obey the axioms of the probability calculus, there is a series of bets, each of which your credences require you to take, but which, when added together, lose you money for sure. In the traditional Dutch Book argument for Conditionalization, we show that, if you plan to update your prior in some way other than by Conditionalization, then there is a series of bets, each of which your priors require you to take, and, for each of your possible posteriors after following the updating plan, there is a series of bets, each of which those posteriors require you to take, but which, when added together, lose you money for sure. Rescorla's improved Dutch

Book argument for Conditionalization does not require a different series of bets for each possible posterior; rather, there is a series of bets, each of which your priors require you to take, and there is a series of bets, each of which *all your possible posteriors* require you to take, but which, when added together, lose you money for sure. The same is true of the Dutch Book arguments for Weak and Strong GRP.

First, let's specify two different ways in which a prior credence function and a set of possible posterior credence functions, taken together, might be flawed; the Dutch Book arguments will take some subset of these flaws to indicate irrationality. Instead of talking about bets specifically, we'll talk generally about any sort of decision-theoretic act. Some terminology:

- A *possible world* is a classically consistent assignment of truth values to the propositions on which the individual's credence function is defined—that is, the propositions in \mathcal{F} . We denote the set of possible worlds \mathcal{W} . Since \mathcal{F} is a finite algebra, \mathcal{W} is also finite. What's more, for each possible world w in \mathcal{W} , there is a unique proposition in \mathcal{F} that is true at that world w and at no other. We abuse notation and write w also for that proposition.
- An *act* is a function that assigns to each possible world the amount of utility you receive in that world (where all utility in what follows is measured on the same scale). Throughout, we'll assume that there is some commodity such that your utility is linear in that commodity, and the outcomes of the acts we discuss involve you receiving certain amounts of that commodity. Standard Dutch Book arguments make a similar assumption, and usually take money as a placeholder for this quantity. This assumption ensures two things: (i) there are constant acts—acts which give you the same amount of the commodity at each world; (ii) the utility of choosing one act in one decision problem and another act in a different decision problem is the sum of the individual utilities of the two acts.
- If A is an act and w is a possible world, $A(w)$ is the amount of utility you receive if you choose A at w .
- If m is a utility level, then \bar{m} is the constant act that takes that utility at every possible world; that is, $\bar{m}(w) = m$, for all worlds w .
- A *decision problem* is a set of available acts.
- If act A belongs to one decision problem and act A' belongs to another, then $A + A'$ is the act of choosing A from the first and A' from the second and its utility at w is $(A + A')(w)$.
- If c is a probabilistic credence function c and A and B are acts, c *prefers* A to B if the expected utility of A relative to c is strictly greater than

the expected utility of B relative to c . That is,

$$\sum_w c(w)A(w) > \sum_w c(w)B(w)$$

- If A and B are acts,
 - A *strongly dominates* B if $A(w) > B(w)$ for all possible worlds w ;
 - A *weakly dominates* B if (i) $A(w) \geq B(w)$ for all possible worlds w , and (ii) $A(w) > B(w)$ for some possible worlds w .

Suppose c is your credence function at t and R is the set of credence functions you might have at t' .

Weak Invariable Dutch Strategy (c, R) is vulnerable to a *weak invariable Dutch strategy* if there are acts A, B, A' , and B' such that

- (i) c prefers A to B ;
- (ii) each c_i in R prefers A' to B'
- (iii) $B + B'$ weakly dominates $A + A'$.

Weak Variable Dutch Strategy (c, R) is vulnerable to a *weak variable Dutch strategy* if there are acts $A, B, A'_1, B'_1, \dots, A'_n, B'_n$ such that

- (i) c prefers A to B ;
- (ii) each c_i in R prefers A'_i to B'_i
- (iii) $B + B'_i$ weakly dominates $A + A'_i$, for all $i = 1, \dots, n$.

Strong Dutch Strategy (c, R) is vulnerable to a *strong Dutch strategy* if there are acts A, B, A' , and B' such that

- (i) c prefers A to B ;
- (ii) each c_i in R prefers A' to B'
- (iii) $B + B'$ strongly dominates $A + A'$.

We can now state our Dutch Book theorem for Weak and Strong GRP:

Theorem 1

- (I) (a) If (c, R) violates Weak GRP, it is vulnerable to a strong Dutch strategy.
 (b) If (c, R) satisfies Weak GRP, it is not vulnerable to a strong Dutch strategy.
- (II) (a) If (c, R) violates Strong GRP, it is vulnerable to a weak invariable Dutch strategy.

(b) If (c, R) satisfies Strong GRP, it is not vulnerable to a weak invariable or weak variable Dutch strategy.

I will now set out of the proof of Theorem 1(I)(a). As so often in the theorems that accompany Dutch Book arguments, it relies on the Separating Hyperplane Theorem.⁷ The proof of Theorem 1(II)(a) is very similar, but appeals to the Supporting Hyperplane Theorem instead. In this context, the Separating Hyperplane Theorem says this: Suppose c is a credence function that is not a convex combination of the credence functions in R , as is the case when (c, R) violates Weak GRP. Then there is an act A and two real numbers, m and m' , such that, for any c_i in R ,

$$\sum_{w \in \mathcal{W}} c_i(w)A(w) < m < m' < \sum_{w \in \mathcal{W}} c(w)A(w)$$

So

- $m' < \sum_{w \in \mathcal{W}} c(w)A(w)$; and
- $-m < \sum_{w \in \mathcal{W}} c_i(w)(-A(w))$, for $i = 1, \dots, n$.

Then c prefers A to the constant act $\overline{m'}$, while each c_i prefers $-A$ to the constant act $\overline{-m}$. But

$$(A + (-A))(w) = 0 < m' - m = (\overline{m'} + \overline{-m})(w)$$

That is, $\overline{m'} + \overline{-m}$ strongly dominates $A + (-A)$. Thus, (c, R) is vulnerable to a strong Dutch strategy, as required.

Next, we prove Theorem 1(II)(b). The proof of Theorem 1(I)(b) is very similar. Suppose (c, R) does satisfy Weak GRP. Then $c = \sum_i \lambda_i c_i$. Now, suppose c prefers A to B and c_i prefers A'_i to B'_i . Then

- $\sum_w c(w)B(w) < \sum_w c(w)A(w)$
- $\sum_w c_i(w)B'_i(w) < \sum_w c_i(w)A'_i(w)$, for $i = 1, \dots, n$.

⁷In general, the Separating Hyperplane Theorem says this: Suppose we have two subsets X and Y of \mathbb{R}^n that are disjoint, non-empty, and convex. Then there is a vector z and a real number r such that, for all x in X and y in Y ,

$$x \cdot z = \sum_{i=1}^n x_i z_i < r < \sum_{i=1}^n y_i z_i = y \cdot z$$

Then

$$\begin{aligned}
& \sum_{w,i} \lambda_i c_i(w) (B + B'_i)(w) \\
= & \sum_{w,i} \lambda_i c_i(w) B(w) + \sum_{w,i} \lambda_i c_i(w) B'_i(w) \\
= & \sum_w c(w) B(w) + \sum_{w,i} \lambda_i c_i(w) B'_i(w) \quad \text{since } c = \sum_i \lambda_i c_i \\
= & \sum_w c(w) B(w) + \sum_i \lambda_i \sum_w c_i(w) B'_i(w) \\
< & \sum_w c(w) A(w) + \sum_i \lambda_i \sum_w c_i(w) A'_i(w) \\
= & \sum_w c(w) A(w) + \sum_{w,i} \lambda_i c_i(w) A'_i(w) \\
= & \sum_{w,i} \lambda_i c_i(w) A(w) + \sum_{w,i} \lambda_i c_i(w) A'_i(w) \quad \text{since } c = \sum_i \lambda_i c_i \\
= & \sum_{w,i} \lambda_i c_i(w) (A + A'_i)(w)
\end{aligned}$$

So it cannot be that $(A + A'_i)(w) < (B + B'_i)(w)$ for all w . □

4 The Accuracy Dominance argument for GRP

In the Accuracy Dominance argument for Probabilism, we show that, if your credence function does not obey the axioms of the probability calculus, there is an alternative credence function, defined on exactly the same propositions, that has greater accuracy than yours for sure. In Briggs and Pettigrew's Accuracy Dominance argument for Conditionalization, we show that, if you plan to update your prior in some way other than by Conditionalization, then there is an alternative prior and an alternative updating plan that have greater total accuracy than your prior and updating plan for sure. In the Accuracy Dominance argument for Weak GRP, we show that, if your prior is not a convex combination of your possible posteriors, there is an alternative prior and alternative possible posteriors, each paired with one of your possible posteriors, such that if you were to replace your prior with the alternative prior and each of the posteriors with its paired alternative, you'd have greater total accuracy for sure.

First, a couple of quick words about measuring accuracy. Briggs and Pettigrew's argument, like the increasingly standard Accuracy Dominance argument for Probabilism, assumes that our measures of inaccuracy have three properties (Predd et al., 2009):

Additivity The inaccuracy of a whole credence function is the sum of the inaccuracy of the credences it assigns.

More precisely: If \mathfrak{I} is a legitimate measure of the inaccuracy of a credence function at a world, then there is, for each X in \mathcal{F} , a scoring rule $\mathfrak{s}_X : \{0, 1\} \times [0, 1] \rightarrow [0, \infty]$ such that, for any credence function c defined on \mathcal{F} and any world w ,

$$\mathfrak{I}(c, w) = \sum_{X \in \mathcal{F}} \mathfrak{s}_X(v_w(X), c(X))$$

where $v_w(X) = 0$ if X is false at w and $v_w(X) = 1$ if X is true at w . In this case, we say that \mathfrak{s} *generates* \mathfrak{I} .

Continuity The inaccuracy of a credence is a continuous function of that credence.

More precisely: If \mathfrak{I} is a legitimate measure of inaccuracy that is generated by the scoring rule \mathfrak{s} , then, for all X in \mathcal{F} , $\mathfrak{s}_X(1, x)$ and $\mathfrak{s}_X(0, x)$ are continuous functions of x .

Strict Propriety Each credence expects itself to be most accurate.

More precisely: If \mathfrak{I} is a legitimate measure of inaccuracy that is generated by the scoring rule \mathfrak{s} , then, for all X in \mathcal{F} and $0 \leq p \leq 1$,

$$p\mathfrak{s}_X(1, x) + (1 - p)\mathfrak{s}_X(0, x)$$

is uniquely minimized, as a function of x , at $x = p$.

When \mathfrak{I} satisfies these three properties, we say that it is an *additive and continuous strictly proper inaccuracy measure*.

Next, let's specify a flaw that a prior and set of possible posteriors might jointly have.

Accuracy Domination A pair $(c, R = \{c_1, \dots, c_n\})$ consisting of a prior c and a set R of possible posteriors is *accuracy dominated* iff, for all legitimate inaccuracy measures \mathfrak{I} , there is an alternative pair $(c^*, R^* = \{c_1^*, \dots, c_n^*\})$ such that, for all possible worlds w and all $i = 1, \dots, n$,

$$\mathfrak{I}(c^*, w) + \mathfrak{I}(c_i^*, w) < \mathfrak{I}(c, w) + \mathfrak{I}(c_i, w)$$

And now our theorem:

Theorem 2 (c, R) *violates Weak GRP iff it is accuracy dominated.*

Proof of Theorem 2. Given a set of possible posteriors, $R = \{c_1, \dots, c_n\}$, define the following set of $(n + 1)$ -dimensional vectors of credence functions:

$$\bar{R} = \{(v_w, c_1, \dots, c_{i-1}, v_w, c_{i+1}, \dots, c_n) : w \in \mathcal{W}, i = 1, \dots, n\}$$

where, again, v_w is the credence function with $v_w(X) = 0$ if X is false at w and $v_w(X) = 1$ if X is true at w . Then:

Lemma 3 *If c violates Weak GRP, then*

$$(c, c_1, \dots, c_n) \notin \bar{R}^+$$

where \bar{R}^+ is the convex hull of \bar{R} .⁸

Proof of Lemma 3. To prove this, we prove the contrapositive. Suppose $(c, c_1, \dots, c_n) \in \bar{R}^+$. Then there are $0 \leq \lambda_{w,i} \leq 1$ such that $\sum_w \sum_i \lambda_{w,i} = 1$ and

$$(c, c_1, \dots, c_n) = \sum_w \sum_i \lambda_{w,i} (v_w, c_1, \dots, c_{i-1}, v_w, c_{i+1}, \dots, c_n)$$

So,

$$c = \sum_w \sum_i \lambda_{w,i} v_w$$

and

$$c_j = \sum_w \lambda_{w,j} v_w + \sum_w \sum_{i \neq j} \lambda_{w,i} c_j$$

So

$$\left(\sum_w \lambda_{w,j} \right) c_j = \sum_w \lambda_{w,j} v_w$$

So let $\lambda_j = \sum_w \lambda_{w,j}$. Then, for $j = 1, \dots, n$,

$$\lambda_j c_j = \sum_w \lambda_{w,j} v_w$$

And thus

$$\sum_i \lambda_i c_i = \sum_i \sum_w \lambda_{w,i} v_w = c$$

Thus, c is in \bar{R}^+ , and c satisfies Weak GRP.

Return to Proof of Theorem 2. Now, we appeal to two central facts about additive and continuous strictly proper inaccuracy measures:

Lemma 4 (Proposition 2, (Predd et al., 2009)) *Suppose \mathfrak{J} is an additive and continuous strictly proper inaccuracy measure. Then there is a Bregman divergence \mathfrak{D} such that, for any credence function c and any world w ,*⁹

$$\mathfrak{J}(c, w) = \mathfrak{D}(v_w, c)$$

Thus, the inaccuracy of c at w is the divergence from v_w to c .

⁸For the definition of the convex hull of a set, see footnote 6.

⁹A Bregman divergence is a certain sort of function that takes pairs of credence functions defined on \mathcal{F} and returns a non-negative real number or infinity. It might be thought of as measuring the distance from one credence function to the other, but it is not required to have certain basic properties of a measure of distance, such as symmetry and the triangle inequality. For our purposes here, we don't need to know anything about them except what is stated in Lemmas 4 and 5.

Lemma 5 (Proposition 3, (Predd et al., 2009)) Suppose \mathfrak{D} is a Bregman divergence and \mathcal{P} is a set of credence functions. Then, if c is not in \mathcal{P}^+ , then there is c^* in \mathcal{P}^+ such that, for all p in \mathcal{P} ,¹⁰

$$\mathfrak{D}(p, c^*) < \mathfrak{D}(p, c)$$

Now, suppose \mathfrak{J} is an additive and continuous strictly proper inaccuracy measure and \mathfrak{D} is its accompanying Bregman divergence. Now, suppose (c, R) violates Weak GRP. Then

$$(c, c_1, \dots, c_n) \notin \overline{R}^+$$

So, by Lemma 5, there is

$$(c^*, c_1^*, \dots, c_n^*) \in \overline{R}^+$$

such that, for all w and i ,

$$\begin{aligned} \mathfrak{D}(v_w, c^*) + \mathfrak{D}(c_1, c_1^*) + \dots + \mathfrak{D}(c_{i-1}, c_{i-1}^*) + \\ \mathfrak{D}(v_w, c_i^*) + \mathfrak{D}(c_{i+1}, c_{i+1}^*) + \dots + \mathfrak{D}(c_n, c_n^*) < \\ \mathfrak{D}(v_w, c) + \mathfrak{D}(c_1, c_1) + \dots + \mathfrak{D}(c_{i-1}, c_{i-1}) + \\ \mathfrak{D}(v_w, c_i) + \mathfrak{D}(c_{i+1}, c_{i+1}) + \dots + \mathfrak{D}(c_n, c_n) \end{aligned}$$

Now, $\mathfrak{D}(c_i, c_i^*) \geq 0$ for all $i = 1, \dots, n$. Furthermore, $\mathfrak{D}(c_i, c_i) = 0$. So we can infer:

$$\mathfrak{D}(v_w, c^*) + \mathfrak{D}(v_w, c_i^*) < \mathfrak{D}(v_w, c) + \mathfrak{D}(v_w, c_i)$$

And so, by Lemma 4,

$$\mathfrak{J}(c^*, w) + \mathfrak{J}(c_i^*, w) < \mathfrak{J}(c, w) + \mathfrak{J}(c_i, w)$$

as required.

Next, suppose (c, R) does satisfy Weak GRP. Then

$$c = \sum_i \lambda_i c_i$$

Now consider c^*, c_1^*, \dots, c_n^* . Then, since \mathfrak{J} is strictly proper:

- $\sum_w c(w) \mathfrak{J}(c, w) \leq \sum_w c(w) \mathfrak{J}(c^*, w)$
- $\sum_w c_i(w) \mathfrak{J}(c_i, w) \leq \sum_w c(w) \mathfrak{J}(c_i^*, w)$, for $i = 1, \dots, n$.

¹⁰Recall: \mathcal{P}^+ is the convex hull of \mathcal{P} .

Then

$$\begin{aligned}
& \sum_{w,i} \lambda_i c_i(w) (\mathfrak{J}(c, w) + \mathfrak{J}(c_i, w)) \\
= & \sum_{w,i} \lambda_i c_i(w) \mathfrak{J}(c, w) + \sum_{w,i} \lambda_i c_i(w) \mathfrak{J}(c_i, w) \\
= & \sum_w c(w) \mathfrak{J}(c, w) + \sum_{w,i} \lambda_i c_i(w) \mathfrak{J}(c_i, w) \quad \text{since } c = \sum_i \lambda_i c_i \\
= & \sum_w c(w) \mathfrak{J}(c, w) + \sum_i \lambda_i \sum_w c_i(w) \mathfrak{J}(c_i, w) \\
\leq & \sum_w c(w) \mathfrak{J}(c^*, w) + \sum_i \lambda_i \sum_w c_i(w) \mathfrak{J}(c_i^*, w) \\
= & \sum_w c(w) \mathfrak{J}(c^*, w) + \sum_{w,i} \lambda_i c_i(w) \mathfrak{J}(c_i^*, w) \\
= & \sum_{w,i} \lambda_i c_i(w) \mathfrak{J}(c^*, w) + \sum_{w,i} \lambda_i c_i(w) \mathfrak{J}(c_i^*, w) \quad \text{since } c = \sum_i \lambda_i c_i \\
= & \sum_{w,i} \lambda_i c_i(w) (\mathfrak{J}(c^*, w) + \mathfrak{J}(c_i^*, w))
\end{aligned}$$

So it cannot be that $\mathfrak{J}(c^*, w) + \mathfrak{J}(c_i^*, w) < \mathfrak{J}(c, w) + \mathfrak{J}(c_i, w)$ for all w in \mathcal{W} .
 \square

5 Problems with Reflection

We are left, then, with two arguments for Weak GRP (and one for Strong GRP). If you violate Weak GRP, your prior and possible posteriors are jointly flawed in two ways. First, they require you to choose a dominated pair of acts in certain situations. Second, there are alternatives that accuracy dominate you. Provided you agree that those flaws render your priors and possible posteriors jointly irrational, we have two arguments for Weak GRP, one pragmatic, one epistemic. And, as we noted above, with Weak GRP in place, we can derive Conditionalization without assuming that whatever we become certain of is true.

However, a natural worry arises. There are countless putative counterexamples to van Fraassen's original formulation of the Reflection Principle (Talbot, 1991; Christensen, 1991; Arntzenius, 2003; Briggs, 2009). Do they not tell equally against the formulations we've given here? I think not, though as I'll explain I'm less than fully certain.

On van Fraassen's original formulation, the principle governs only your prior credence function. It imposes no constraints on your future credences, but only on the relationship between your prior opinions about your future credences in certain matters and your prior credences in those matters. It says that your current credence in X , conditional on your future credence in X being r , should be r .

Let's consider one of the counterexamples due to David Christensen (1991). I know that, between t and t' , I will learn no new evidence, but I will begin to feel the effects of a hallucinogenic drug that I've just taken. I also know that one of the effects is that it will make me very confident that I can fly. The Reflection Principle deems irrational a credence function that is simultaneously certain that my future credence that I can fly will be high, and very confident that I cannot fly. But that seems a perfectly rational response to the evidence I currently have.

What does the Weak GRP say about this case? It says that, if my only possible future credence function is very confident that I can fly, while my current credence function is very confident I cannot, then they are jointly irrational. But that isn't a problem. After all, I'd also agree that they are jointly irrational, because I judge my future credences to be irrational—they're the result of taking a hallucinogen, and my aerial competence is not one of the truths typically revealed by such drugs. The upshot: this counterexample doesn't tell equally against Weak GRP. The reason is that the original Reflection Principle is a purely synchronic norm, while Weak GRP is genuinely diachronic.

This might lead you to wonder how Weak GRP can help to establish Conditionalization, which I noted at the beginning is a synchronic norm governing the relationship between your credence function at t and your updating plan at t , not between your credence function at t and your possible credence functions at t' . The reason is that an updating plan is a commitment to having certain credence functions at t' given certain learning experiences prior to that. So it is reasonable to judge it by looking at what would happen were you to follow that plan and update accordingly. If you were to do that, then you'd have certain possible posteriors. And if your prior together with those possible posteriors are irrational because they violate Weak GRP, then your prior together with the updating plan that demands those posteriors is similarly irrational.

However, a different sort of counterexample to reflection principles does concern me. Suppose you are a permissivist about rationality. You have a certain body of evidence that permits two different credence functions c and c' . In the evening, at time t , your credence function is c —perhaps in the evening, after the stress of a long day, you're more inclined to the more extreme of the two permissible responses. But, from past experience, you know that, by the morning, after a long sleep during which you will gain no new evidence, your credence function will be c' —perhaps you know that, after some rest, the stress will dissipate somewhat and you'll instead favour the less extreme permissible response. So c' is your only possible future credence function. But it is not c . So, together, they violate Weak GRP.

Unlike the case of the hallucinogenic drug above, in this case both credence functions are rationally permissible. So, if they are jointly irrational,

their individual irrationality cannot be the reason. Instead, it must lie in the move from one to the other. I find myself without a clear sense of whether Weak GRP gives the correct verdict in this case. In the drug case, I had a strong conviction that the prior credence function that violates Reflection is perfectly rational. So in that case Reflection was in direct conflict with my judgments. I have no strong conviction that the prior c and posterior c' are jointly rational in this case. So Weak GRP does not conflict with my judgments. But equally I have no strong conviction that they are jointly irrational. So Weak GRP does add something to my pre-theoretic judgments. It makes a call in a case where I don't. Ultimately, I must leave it to the reader to determine whether they think this verdict sinks Weak GRP and with it the Dutch Book and Accuracy Dominance arguments in its favour, as well as the argument for Conditionalization that follows after it.

References

- Arntzenius, F. (2003). Some Problems for Conditionalization and Reflection. *Journal of Philosophy*, 100(356-70).
- Briggs, R. A. (2009). Distorted Reflection. *Philosophical Review*, 118(1), 59–85.
- Briggs, R. A., & Pettigrew, R. (2018). An accuracy-dominance argument for conditionalization. *Noûs*.
- Brown, P. M. (1976). Conditionalization and expected utility. *Philosophy of Science*, 43(3), 415–419.
- Christensen, D. (1991). Clever Bookies and Coherent Beliefs. *Philosophical Review*, 100(2), 229–247.
- Greaves, H., & Wallace, D. (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459), 607–632.
- Huttegger, S. M. (2013). In Defense of Reflection. *Philosophy of Science*, 80(3), 413–433.
- Jeffrey, R. (1965). *The Logic of Decision*. New York: McGraw-Hill.
- Lewis, D. (1999). Why Conditionalize? In *Papers in Metaphysics and Epistemology*, (pp. 403–407). Cambridge, UK: Cambridge University Press.
- Oddie, G. (1997). Conditionalization, Cogency, and Cognitive Value. *British Journal for the Philosophy of Science*, 48, 533–41.
- Predd, J., Seiringer, R., Lieb, E. H., Osherson, D., Poor, V., & Kulkarni, S. (2009). Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions of Information Theory*, 55(10), 4786–4792.

- Rescorla, M. (2020). An Improved Dutch Book Theorem for Conditionalization. *Erkenntnis*.
- Skyrms, B. (1997). The Structure of Radical Probabilism. *Erkenntnis*, 45, 285–297.
- Talbott, W. J. (1991). Two Principles of Bayesian Epistemology. *Philosophical Studies*, 62(2), 135–150.
- van Fraassen, B. C. (1981). A Problem for Relative Information Minimizers. *The British Journal for the Philosophy of Science*, 32(4), 375–379.
- van Fraassen, B. C. (1995). Belief and the Problem of Ulysses and the Sirens. *Philosophical Studies*, 77(1), 7–37.
- van Fraassen, B. C. (1999). Conditionalization, A New Argument For. *Topoi*, 18(2), 93–96.
- Weisberg, J. (2007). Conditionalization, Reflection, and Self-Knowledge. *Philosophical Studies*, 135(2), 179–97.