

Article

Maximum Relative Entropy Updating and the Value of Learning

Patryk Dziurosz-Serafinowicz

Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, Groningen, 9712 GL, The Netherlands; E-Mail: p.dziurosz-serafinowicz@rug.nl; Tel.: +31-50-363-6161

Academic Editors: Juergen Landes and Jon Williamson

Received: 7 January 2015 / Accepted: 4 March 2015 / Published: 11 March 2015

Abstract: We examine the possibility of justifying the principle of maximum relative entropy (MRE) considered as an updating rule by looking at the value of learning theorem established in classical decision theory. This theorem captures an intuitive requirement for learning: learning should lead to new degrees of belief that are expected to be helpful and never harmful in making decisions. We call this requirement the value of learning. We consider the extent to which learning rules by MRE could satisfy this requirement and so could be a rational means for pursuing practical goals. First, by representing MRE updating as a conditioning model, we show that MRE satisfies the value of learning in cases where learning prompts a complete redistribution of one's degrees of belief over a partition of propositions. Second, we show that the value of learning may not be generally satisfied by MRE updates in cases of updating on a change in one's conditional degrees of belief. We explain that this is so because, contrary to what the value of learning requires, one's prior degrees of belief might not be equal to the expectation of one's posterior degrees of belief. This, in turn, points towards a more general moral: that the justification of MRE updating in terms of the value of learning may be sensitive to the context of a given learning experience. Moreover, this lends support to the idea that MRE is not a universal nor mechanical updating rule, but rather a rule whose application and justification may be context-sensitive.

Keywords: maximum relative entropy; probabilistic updating; the value of learning theorem; decision theory; Skyrms's condition M; the Judy Benjamin problem; context-sensitivity

1. Introduction

Let the probability functions P and P' represent, respectively, an agent's prior and posterior degrees-of-belief functions over an algebra of propositions \mathcal{F} generated by a set of possible worlds W . A rule for changing the agent's prior degrees-of-belief function P over \mathcal{F} in light of new evidence (hereafter, an updating rule) aims to provide an answer to the following problem: given P and some constraint χ imposed on P' , which P' should the agent choose from the set of her posterior degrees-of-belief functions that satisfy χ ? A given constraint χ imposed on P' is supposed to represent a learning experience, and we associate with every learning experience a set \mathcal{P}_χ of posterior degrees-of-belief functions singled out by χ , *i.e.*, $\mathcal{P}_\chi = \{P' : P' \text{ satisfies } \chi\}$. We take it that \mathcal{P}_χ is a closed convex set, *i.e.*, it is determined by a constraint χ , such that if P'_1 and P'_2 satisfy χ , then also, any convex combination of them, $\lambda P'_1 + (1 - \lambda)P'_2$ with $\lambda \in [0, 1]$, will satisfy χ . This type of constraint is called affine.

An updating rule that is subject to considerable discussion among philosophers is the principle of maximum relative entropy (MRE), also known as the rule of minimizing cross-entropy, the principle of minimum discrimination information or Kullback–Leibler divergence. It says that, given P , the partition $\{S_i\}$ of minimal elements in \mathcal{F} and some constraint χ on P' , the agent should choose P' , so as to satisfy χ , while minimizing the relative entropy with respect to P as measured by the following function:

$$\text{RE}(P, P') = \sum_i P'(S_i) \log \frac{P'(S_i)}{P(S_i)}$$

That is, by MRE, an updater should adopt as her posterior degrees-of-belief function, from those defined over \mathcal{F} and satisfying χ , the one that is RE-closest to her prior degrees-of-belief function defined over \mathcal{F} . RE; thus, it can be seen as measuring the “distance” between P and the possible P' 's that satisfy χ . Additionally, $\text{RE} = 0$ just in case $P = P'$. Of course, RE is not a distance measure in the mathematical sense, for it is not symmetric.

Much of the controversy surrounding MRE concerns its status. At least four main views on this issue can be distinguished. According to the first view [1], MRE is a generally valid rule of updating one's degrees of belief from which the two well-known conditionalization rules, to wit, Bayes's rule and Jeffrey's rule, derive their normative force. The second view denies the very idea of MRE's universal validity. Within this camp, some [2–4] argue that in certain situations, it conflicts with Bayes's rule; others [5,6] argue that it leads to counterintuitive consequences in the Judy Benjamin case, which is a case of updating on a conditional proposition; and some [7,8] argue, quite generally, that MRE is just one of many updating rules and, as such, is applicable in the right circumstances. In the third view [9], MRE can be regarded, under certain conditions, as a special case of Bayes's rule. Finally, in the fourth view [10], MRE is not a rule for updating one's degrees of belief, but rather a rule for statistical supposing. These views have their merits, although none have achieved widespread acceptance.

However, there is yet another foundational question concerning MRE, a question that might be posed independently of the aforementioned concern. This is the question of whether, and if so, how, MRE can be justified as a method of updating one's degrees of belief. Surprisingly, there have been relatively few attempts to answer this question. The most notable among them are Shore and Johnson's [11,12] justification by consistency and Grünwald's [13] minimax decision-theoretic justification. In contrast,

there are several existing justifications of the two most prominent updating rules, to wit, Bayes's rule and Jeffrey's rule. Bayes's rule is justified on the grounds that it is both a pragmatically and epistemically rational way of updating. The pragmatic rationality of this rule is established by a diachronic Dutch book argument, which shows that if you update your degrees of beliefs other than by Bayes's rule, then you are susceptible to a collection of bets ensuring a negative net pay-off, come what may. Various accuracy-based arguments show that Bayes' rule is also epistemically rational. In particular, they show that Bayesian updating minimizes the expected inaccuracy [14]. Similarly, various Dutch book arguments support Jeffrey's rule, establishing its pragmatic rationality.

The aim of this paper is to examine the possibility of justifying MRE updating by linking it to the value of learning theorem introduced to the philosophical literature by Savage [15] and Good [16]. The value of learning theorem may be viewed as capturing an intuitive requirement of rationality for learning. The requirement says that learning should lead to new degrees of belief that are expected to be helpful and never harmful in making decisions. Call this requirement the value of learning. The notion of rationality that it alludes to is essentially pragmatic: we consider whether an opinion shift ruled by MRE is rational for an agent who always chooses that act that maximizes her expected utility. However, as recently argued in [17], we can also think of the value of learning as a necessary requirement for one's opinion shift to count as genuine learning. Of course, in this view, there might be other features of genuine learning that are not captured by the value of learning. Therefore, it might not be a sufficient condition. Importantly, it has been shown that the value of learning holds for both Bayes's rule [16] and Jeffrey's rule [18].

We show that updating by MRE satisfies the value of learning in cases where the constraint reporting one's learning experience concerns a complete redistribution of one's degrees of belief over a partition of propositions. Our strategy will be to exploit a link between a particular generalized model of Bayesian conditioning and updating by MRE on a partition of propositions. The generalized model of conditioning allows us to assign second-order degrees of belief to propositions about first-order ones and to condition the former on propositions concerning the latter. If we interpret the second-order degrees of belief as one's priors and the first-order ones as one's posteriors, then we can condition prior degrees of belief on propositions about the posterior ones. In this set-up, we can represent, under certain conditions, updating by MRE on a partition as a form of conditioning on a proposition specifying posterior degrees of belief for each member of that partition. However, there are other types of constraints to which MRE updating can be applied. In particular, these might involve a constraint to the effect that one should assign a conditional posterior degree of belief for some proposition, given another proposition. We show that whether or not MRE updating leads to the value of learning theorem in response to such a constraint crucially depends on how broadly the constraint is described. If this constraint can be described effectively as a complete redistribution of one's degrees of belief over a partition of propositions, the value of learning theorem holds. However, if it cannot be so formulated, then the value of learning theorem cannot be established. We explain why this is so: contrary to what the value of learning theorem requires, in such cases, the MRE updater's prior degrees of belief are not equal to the expectation of her possible posterior degrees of belief.

There is yet another angle from which we might look at the main result of this paper. It is often said that MRE is an updating rule that prescribes modesty or minimal revision for the agent's opinion

shifts. As characterized in [5] (p. 376), MRE is “the rule that one should not jump to unwarranted conclusions, or add capricious assumptions, when accommodating one’s belief state to the deliverances of experience”. Minimizing RE under some constraints imposed on posterior degrees of belief is a way, but by no means the only way, to make the idea of modesty more precise: the agent adopts the posterior degree-of-belief function that meets the constraints reporting her learning experience and is RE-closest to her prior degree-of-belief function. Under this procedure, the existence of a uniquely maximally modest P' satisfying a given constraint is guaranteed, since \mathcal{P}_χ is a closed convex set. However, why should we value such modest opinion shifts? Of course, modesty might itself be a virtue that does not require further justification. Be that as it may, modesty might also be viewed as a rational tool for pursuing other goals. What this paper shows is that it is not always true that revising degrees of belief by dint of MRE leads to modest new degrees of belief that are expected to be helpful and never harmful for one’s decisions.

2. The Value of Learning and Bayes’s Rule

It is rather uncontroversial to say that a change in one’s degrees of belief may bring consequences for one’s decisions. Suppose that you have to decide now whether to act on the basis of your current information or to perform a cost-free experiment to obtain further information, update your degrees of belief and then act. For example, you have to decide whether to submit your paper to a journal now or to pursue some line of research, update your degrees of belief about the content of your paper and then decide whether to submit it. What should you do?

There is a striking result in decision theory, due originally to Ramsey and revived by Savage [15] and Good [16], that gives an answer to the aforementioned concern. Informally put, the theorem states that the prior expectation of making an informed decision is at least as great as the expected utility of making an uninformed decision and is strictly greater if it is not the case that the maximum expected utility of an act is the same for all possible experimental results (or equivalently, if at least one of the experimental results could alter the choice of one’s actions). This theorem is known in the literature as the value of learning or the value of knowledge theorem.

In its original form, the theorem has been proven in the context of Bayes’s rule of conditioning. As shown by Good, Bayes’s rule implies the value of learning theorem. To present Good’s argument, let us introduce the following assumptions:

- Let $\mathcal{A} = \{A_1, \dots, A_m\}$ be a finite set of actions and $\mathcal{S} = \{S_1, \dots, S_n\}$ a finite set of states of the world.
- For each combination of A_i and S_j , we introduce a utility function $U(A_i \wedge S_j)$.
- Let $\mathcal{E} = \{E_1, \dots, E_k\}$ be a finite partition of experimental outcomes.
- Assume that the agent is an expected utility maximizer, that is she chooses the act A_i that maximizes her expected utility given by:

$$\sum_j P(S_j)U(A_i \wedge S_j),$$

where $P(S_j)$ is the agent’s prior degree of belief in S_j .

- The agent’s learning experience is reported by the constraint χ saying that one should assign posterior degree of belief one to some E_k . Then, the associated set of posterior degrees-of-belief functions is $\mathcal{P}_\chi = \{P' : P'(E_k) = 1\}$. Bayes’s rule prescribes you to choose from that set the posterior degrees-of-belief function P' that satisfies the constraint and is defined as follows:

(Bayes’s rule) For all j ,

$$P'(S_j) = P(S_j|E_k),$$

provided that $P(E_k) > 0$.

That is, your posterior degree of belief in S_j equals your prior degree of belief in S_j conditional on E_k .

- The experiment is costless.

For simplicity’s sake, we consider only finite sets of states. The value of learning theorem carries over to infinite sets of states if the degrees-of-belief function is countably additive.

Suppose that the agent is faced with the following decision problem. She has to decide whether to act now or to wait until the experiment is performed, update her degrees of belief by Bayes’s rule and then act. Since the agent is an expected utility maximizer, the present value of her deciding now, without performing the experiment, is:

$$\begin{aligned} \max_i \sum_j P(S_j)U(A_i \wedge S_j) &= \max_i \sum_k \sum_j P(S_j|E_k)P(E_k)U(A_i \wedge S_j) \\ &= \max_i \sum_k \sum_j \frac{P(E_k|S_j)P(S_j)}{P(E_k)}P(E_k)U(A_i \wedge S_j) \\ &= \max_i \sum_k \sum_j P(S_j)P(E_k|S_j)U(A_i \wedge S_j), \end{aligned} \tag{1}$$

which is the expected value of act A_i with the highest expected utility.

The present value of making an informed decision is given as follows. Suppose that E is the true member of \mathcal{E} . Then, the posterior value of making a decision informed by E is the value of act A_i with the highest expected utility with respect to the conditional degree of belief $P(S_j|E)$:

$$\max_i \sum_j P(S_j|E)U(A_i \wedge S_j). \tag{2}$$

Given (2), the present value of making a decision conditional on E is calculated by:

$$\sum_k P(E_k) \max_i \sum_j P(S_j|E_k)U(A_i \wedge S_j) = \sum_k \max_i \sum_j P(S_j)P(E_k|S_j)U(A_i \wedge S_j), \tag{3}$$

which is the prior expectation of the posterior value of making an informed decision.

Note that Equations (1) and (3) differ only in the order of the \max_i and the \sum_k operations. Additionally, by Jensen’s inequality, for any real-valued and convex function $f(k, i)$ of k and i , $\sum_k \max_i f(k, i) \geq \max_i \sum_k f(k, i)$, with strict inequality if it is not the case that $\max_i f(k, i)$ is the same for all k . Hence, it follows that Equation (3) is at least as great as Equation (1), with strict inequality if it is not the case that the act A_i maximizes the expected utility irrespective of which of the E_k ’s hold true.

The value of learning theorem carries an important philosophical message for someone who evaluates learning and updating rules in terms of their potential consequences for decisions. The message is that, from the perspective of maximizing expected utility, a change in one's degrees of belief could make one's decisions better and never worse. That is, acquiring information by way of an update is expected to be helpful and never harmful. Of course, this result does not hold unconditionally. It rests on a few substantial assumptions. First of all, it is set up in the framework of Savage's decision theory in which states of the world and acts are stochastically independent in the sense that choosing an act does not give you information about which state of the world is true. Likewise, one's decision whether to perform an experiment is stochastically irrelevant to the states of the world. Notice, however, that updating on experimental outcomes may alter your degrees of belief about the states. Second, the states, acts and utilities are the same before and after updating your degrees of belief. Third, it is assumed that you are an expected utility maximizer before and after updating.

It is important to recognize that the agent assesses the value of making an informed decision from her current perspective, without knowing which of the experimental outcomes is true. To assess this value, she takes the expectation of Equation (2) with respect to the unknown E_k . This, in turn, shows how her prior degrees of belief must be related to her possible posterior degrees of belief. Since she knows that she will update by Bayes's rule, it follows that for each j , her prior degree of belief in S_j must be equal to the expectation of her conditional prior degrees of belief, the $P(S_j|E_k)$'s; that is:

$$P(S_j) = \sum_k P(E_k)P(S_j|E_k), \quad (4)$$

where the sum extends over all k , such that $P(E_k) > 0$. This is an elementary observation. However, what happens if Bayes's rule is not assumed? In the next section, we will suggest a more general answer to the question of how the agent's prior degrees-of-belief function should be related to her possible posterior degrees-of-belief functions for the value of learning to be satisfied. This answer involves focusing on Skyrms's condition M.

3. Condition M and the Value of Learning

Does the value of learning imply a particular way in which one's prior and one's possible posterior degrees of belief are related? In this section, we give an affirmative answer to this question by exploring Skyrms's condition M. We will present this condition within the framework of an unstructured and opaque degrees-of-belief change called by Skyrms [19,20] black-box learning. It is unstructured in the sense that we do not know how the agent updates her degrees of belief (*i.e.*, what rule she adopts as her updating policy) and what the constraint that prompts the shift in her degrees of belief is. The only thing we know is the effect of her learning experience on her posterior degrees of belief.

Black-box learning is a generalized model of learning. According to it, an epistemic agent starts with a prior degrees-of-belief function, passes through a black-box learning experience and ends up with a posterior degrees-of-belief function. Thus, the agent only knows the input (prior degree-of-belief function) and the output (posterior degrees-of-belief function). Here, the learning process is not transparent: the agent cannot go into the black-box and see what is inside. In particular, she cannot say whether she learned a proposition with certainty or redistributed her degrees of belief over a partition

of propositions. That is, she cannot specify a constraint that prompts the shift in her degrees of belief. Likewise, she cannot specify a rule of updating that would deal with her learning episode. For example, she does not expect that she would learn a proposition as a result of her interaction with the environment, yet she might think about this experience and revise her opinion on the basis of her thoughts. More precisely, black-box learning may be described as follows. Let an agent's degrees-of-belief space be a triple (W, \mathcal{F}, P) , where W is a set of worlds that the agent considers possible, the elements in \mathcal{F} are propositions about which the agent has an opinion and P is the agent's degree-of-belief function. Suppose that the agent is in a learning situation where she expects her degrees-of-belief function over \mathcal{F} to change from P to one of the posterior degrees-of-belief functions in the set $\{P'\}$, resulting from her interaction with the environment. Since her learning is described only by the effect on her possible posterior degrees-of-belief functions, we can enlarge her degrees-of-belief space by adding the posterior degrees-of-belief function as a random variable. As a result, the agent might have second-order degrees of belief over propositions about the first-order ones. The first-order degrees of belief are her possible posterior degrees of belief. By doing so, we get a higher-order probability structure in the sense proposed in [21]. Such a structure may be represented by (W, \mathcal{F}, P, P') , where \mathcal{F} is an algebra of propositions, subsets from W , P is one's prior degree-of-belief function over \mathcal{F} and P' is a measurable function defined as $P' : \mathcal{F} \times [0, 1] \rightarrow \mathcal{F}$. Let the proposition about one's posterior degrees of belief be denoted by $X_{P'}$. The proposition says that the posterior degree-of-belief function over \mathcal{F} is given by P' .

Could a black-box learner satisfy the value of learning? Recall that the black-box learner has no updating rule at his disposal and no constraint that prompts his degrees-of-belief shift. One might, thus, be suspicious as to whether black-box learning could be even justified. After all, we deal with a situation where one expects one's degrees of belief will change as a result of an interaction with the environment without being confident that the change will be prompted by something learned. Additionally, a black-box learning situation does not exclude the possibility that reasons other than learning might prompt one's degrees of belief change. In particular, one might expect that one's degrees of belief will change by taking a drug that makes one confident that one can fly, by memory loss or by being brainwashed. Skyrms [19] shows convincingly that a sufficient condition for one's degrees-of-belief change in black-box learning to satisfy the value of learning is the following:

(M) An agent's prior degrees-of-belief function ought to be such that, for all j and for any possible posterior degrees-of-belief function P' :

$$P(S_j | X_{P'}) = P'(S_j),$$

providing that $P(X_{P'}) > 0$.

That is, condition M requires one's prior degree of belief in S_j conditional on the proposition about S_j 's posterior degree of belief to be equal to that posterior degree of belief. In [19], M stands for Martingale. A similar principle, known as reflection, has been defended in [22].

Skyrms’ reasoning goes as follows. The agent’s present value of deciding now is the maximum of her prior expectation of posterior expected utility. In symbols,

$$\begin{aligned} \max_i \sum_j P(S_j)U(A_i \wedge S_j) &= \max_i \sum_{P'} \sum_j P(S_j|X_{P'})P(X_{P'})U(A_i \wedge S_j) \\ &= \max_i \sum_{P'} \sum_j P'(S_j)P(X_{P'})U(A_i \wedge S_j). \end{aligned} \tag{5}$$

The posterior value of making a decision informed by $X_{P'}$ is given by:

$$\max_i \sum_j P(S_j|X_{P'})U(A_i \wedge S_j) = \max_i \sum_j P'(S_j)U(A_i \wedge S_j). \tag{6}$$

Given Equation (6), we can calculate the present value of making an informed decision as one’s prior expectation of the value given by Equation (6). That is,

$$\sum_{P'} P(X_{P'}) \max_i \sum_j P'(S_j)U(A_i \wedge S_j). \tag{7}$$

Now, as shown by Skyrms, it is a consequence of Jensen’s inequality that the value given by Equation (7) is at least as great as the value given by Equation (5). Thus, condition M satisfies the value of learning.

What happens if condition M fails? Skyrms [20] shows that if the black-box learner fails to satisfy condition M, then the expected utility of her informed decision could be lower than the expected utility of her uninformed decision. Thus, condition M is both sufficient and necessary for the value of learning to hold. Similarly, Huttegger [17] argues that condition M and the value of learning are in fact equivalent. Assuming Skyrms’s result, Huttegger shows quite generally that if updating one’s degrees of belief satisfies the value of learning, then condition M must hold. Thus, condition M is all we need for the value of learning to hold. To explain the necessity of condition M, suppose that $P(S_j|X_{P'}) = \frac{1}{3}$ and $P'(S_j) = \frac{2}{3}$. Hence, you violate condition M. Consider a bet on S_j conditional on the proposition that $P'(S_j) = \frac{2}{3}$; it costs you \$5 and pays you \$5 if both S_j and the proposition that $P'(S_j) = \frac{2}{3}$ are true. Since you violate condition M, you are vulnerable to a Dutch book, *i.e.*, a set of bets that guarantee you a net loss, come what may. You have to decide now whether to accept this bet or to update your degrees of belief in S_j and then decide. Since your decision to reject this bet now has greater expected utility than your decision to act later and possibly to risk acceptance of this bet, the value of learning theorem fails to hold.

Now, if condition M alone is all that is required for the value of learning to hold, we can determine, by focusing solely on that condition, the way in which one’s prior and posterior degrees of belief should be related for one’s opinion shift to satisfy the value of learning. Additionally, since we deal with a black-box learning situation, this way of relating priors and posteriors must be independent of which updating rule the agent endorses as her updating policy.

It is an immediate consequence of condition M that one’s prior degrees of belief are the expectation of one’s anticipated posterior degrees of belief, *i.e.*, for all j :

$$P(S_j) = \sum_{P'} P'(S_j)P(X_{P'}). \tag{8}$$

In other words, the agent’s prior degree of belief in S_j is a convex combination of her possible posterior degrees of belief in S_j . Given that Equation (8) is a consequence of condition M, if Equation (8)

fails to hold, then condition M cannot be satisfied, and hence, the value of learning theorem cannot be established. Note that Equation (8) does not tell us how the agent arrives at her posterior degrees of belief. After all, Equation (8) characterizes a black-box learner. The basic idea behind Equation (8) is that no matter how the agent arrives at her posterior degrees of belief, her prior degrees of belief are required to be the expectation of her posterior ones.

It is not hard to observe that a Bayesian conditionalizer satisfies Equation (8). If you know that you will update by dint of Bayes's rule, your prior degrees of belief are the expectation of your anticipated posterior ones that are given by the conditional prior degrees of belief. Of course, the important question here is: in what sense one's conditional degrees of belief, the $P(S_j|E_k)$'s, capture one's anticipated degrees of belief that figure in Equation (8). Two interesting answers to this question are given in the literature. First, as pointed out in [23], one might believe with degree one that one will update by Bayes's rule on E_k . Then, one's anticipated future degrees of belief are just the $P(S_j|E_k)$'s. Second, following Easwaran [24], one might view the $P(S_j|E_k)$'s as "plans" to update one's degrees of belief after learning which member of \mathcal{E} is true. Then, the agent's anticipated future degrees of belief are simply her degrees of belief that she plans to have. In my view, both of these answers are plausible ways to find a bridge between one's conditional degrees of belief and one's anticipated future degrees of belief.

In what follows, we show that updating by MRE on a constraint prompting a complete redistribution of degrees of belief over a partition of propositions agrees with a Bayesian model of learning from experience that satisfies Equation (8). This, in turn, leads straightforwardly to the value of learning theorem for MRE. However, we also show that MRE updates on a constraint prompting a change in one's conditional degrees of belief might not lead to the value of learning theorem. We explain that this is because such MRE updates might not coincide with a model of learning that satisfies Equation (8).

4. The Value of Learning and MRE

In general, MRE updating can be applied to a learning situation reported by an affine constraint on posterior degrees of belief. An affine constraint can always be formulated as saying that one's expectation of a random variable, computed relative to one's posterior degrees-of-belief function, has a given value. Examples of such constraint include: (i) a constraint to the effect that she should assign posterior degrees of belief to a partition of propositions without conferring certainty on any of them; or (ii) a constraint to the effect that she should assign a conditional posterior degree of belief for some proposition, given another proposition. For example, to see how Constraint (i) can be expressed as one's expectation of a random variable, suppose that X is a \mathcal{F} -measurable random variable, *i.e.*, a function from W to the real numbers \mathbb{R} . Suppose that the elements of a partition $\{E_1, \dots, E_k\}$ of W are represented as 0,1-valued random variables or indicator functions. The indicator function of E_i , denoted by $I_{E_i}(w)$, can be understood as the truth value of E_i at world w , that is, $I_{E_i}(w) = 1$ if $w \in E_i$, and $I_{E_i}(w) = 0$ otherwise. Since posterior degrees of belief over the members of that partition are equal to the posterior expectations of the indicator functions, (i) may be reformulated as a constraint to the effect that the expectations of these indicator functions, computed with respect to the posterior degrees-of-belief function, get some values in \mathbb{R} . In this section, we show that an MRE update in response to Constraint (i) leads to the value of learning theorem.

To this end, we first introduce the following well-known result. Suppose that the agent's learning experience is reported by the following constraint. Let $\mathcal{E} = \{E_1, \dots, E_k\}$ be a partition of W , and let $q_1, \dots, q_k \in \mathbb{R}^+$ be such that $q_1 + \dots + q_k = 1$. Then, χ is a constraint to the effect that upon learning experience, the agent redistributes her degrees of belief over $\{E_1, \dots, E_k\}$, such that $P'(E_i) = q_i$, for $i = 1, \dots, k$. The agent's set of posterior degrees of belief that satisfy this constraint is given by $\mathcal{P}_\chi = \{P' : P'(E_i) = q_i, i = 1, \dots, k\}$, which is a closed and convex set. Given that the agent updates her degrees of belief by MRE, she chooses from the set \mathcal{P}_χ her posterior degrees-of-belief function that minimizes the distance measured by RE. There is a result showing that if the constraints on posterior degrees of belief concern a whole partition of propositions, RE is uniquely minimized just in case the agent's posterior degrees-of-belief function comes by Jeffrey's rule on the partition $\{E_1, \dots, E_k\}$ (see [1,25]). That is, P' should be such that, for all j :

$$P'(S_j) = \sum_i P(S_j|E_i)q_i \quad (9)$$

That is, P' is a weighted average of the agent's prior conditional degrees of belief for S_j given E_i , for all i , where the weights are the values of posterior degrees of belief for the E_i 's. This result may be summarized by the following proposition:

Proposition 1. *Suppose that $\mathcal{P}_\chi = \{P' : P'(E_i) = q_i, i = 1, \dots, k\}$. Then, $\text{RE}(P, P'') \geq \text{RE}(P, P(\cdot|E_i)q_i)$ for all $P'' \in \mathcal{P}_\chi$, with equality, just in case $P'' = P(\cdot|E_i)q_i$.*

As shown by Jeffrey [26], the agent's posterior degree-of-belief function is equal to the one given by Formula (9) if and only if the following condition holds:

(Rigidity) For all j and all i ,

$$P'(S_j|E_i) = P(S_j|E_i).$$

Rigidity says that the agent's conditional degrees of belief given members of $\{E_1, \dots, E_k\}$ remain intact as she shifts her degrees of belief from P to P' . Since MRE updating on a whole partition $\{E_1, \dots, E_k\}$ is also rigid, there is no surprise that it coincides with Jeffrey's rule. We may look at Rigidity in the case of MRE updating as follows: under RE-minimization, for each member E of $\{E_1, \dots, E_k\}$, the ratios of one's posterior degrees of belief to one's prior degrees of belief about propositions that imply E do not change, *i.e.*, if S_i and S_j , $i \neq j$, imply E , then $\frac{P'(S_j)}{P'(S_i)} = \frac{P(S_j)}{P(S_i)}$.

With this result in hand, we can introduce a way to represent MRE updating in response to Constraint (i) as Bayesian conditioning in an enlarged degrees-of-belief space. This move is mobilized by a general result, due to Diaconis and Zabell [25], when a shift from P to P' in the original smaller space agrees with Bayesian conditioning in some bigger space. A related result, though somewhat different in detail, is defended by Grünwald and Halpern [27]. For a two-element partition of propositions, a similar result is given in [28]. Roughly, the idea is as follows. Suppose that the agent shifts from P to P' by MRE updating on a partition of propositions. Given the agent's learning experience reported by a complete redistribution of her degrees of belief over that partition, we can enlarge the original space by adding the proposition that describes the agent's learning experience and the proposition that describes its absence. The proposition that describes the agent's learning experience is about the values that her posterior

degrees-of-belief function assigns to each member of the partition. Then, under certain conditions, we can show that the MRE update in the original smaller space agrees with Bayesian conditioning in the bigger space.

More precisely, to enlarge the agent’s degrees-of-belief space, we add to the algebra \mathcal{F} a proposition X_{q_i} for each member i of the partition \mathcal{E} . Thus, we require that the underlying space (W, \mathcal{F}) is sufficiently rich. In fact, each element of W specifies a value for q_i , which, in turn, may be regarded as a random variable. X_{q_i} says that the agent’s posterior degree of belief assigned to the i -th member of \mathcal{E} equals q_i . This proposition may be understood as a set of worlds from W at which the posterior degree of belief in E_i equals q_i . Denote the algebra extended by adding such propositions by \mathcal{F}^* . The agent’s prior degree-of-belief function P over \mathcal{F}^* may be viewed as a second-order degree-of-belief function, since it assigns degrees of belief to her degrees of belief that are assigned to the propositions in the smaller original algebra \mathcal{F} . Propositions about which the agent has an opinion and that belong to the extended algebra are the propositions that describe her learning experience reported by Constraint (i), to wit, a learning experience that prompts a complete redistribution over the partition \mathcal{E} . Such propositions specify the agent’s degrees of belief for every member of the partition \mathcal{E} . They may be understood as conjunctions, the $\bigwedge_{i=1}^k X_{q_i}$ ’s, of the X_{q_i} ’s. For convenience, denote such a conjunction by D . Now, if the agent learns such a proposition with certainty, she can Bayes condition in the enlarged algebra. In fact, when she conditions in the enlarged algebra, she assigns second-order degrees of belief to propositions about her first-order ones. Denote such Bayesian conditioning in the enlarged algebra by BC^* . It can be put as follows:

(BC*) For all j and any $D \subseteq W$,

$$P'(S_j) = P(S_j|D),$$

provided that $P(D) > 0$.

The following theorem states that under certain conditions, updating by MRE on a partition \mathcal{E} is representable as BC^* .

Theorem 1. *Suppose that the agent’s prior degrees-of-belief function P obeys the following two conditions:*

- (1) *For all i , $P(E_i|D) = q_i$, provided that $P(D) > 0$.*
- (2) *For all j and all i , $P(S_j|E_i \wedge D) = P(S_j|E_i)$, provided that $P(E_i \wedge D) > 0$.*

Then, for all j , $P(S_j|D) = \sum_i P(S_j|E_i)q_i$.

Proof. Suppose that P satisfies Conditions (1) and (2), $D \subseteq W$ and $E_i \subseteq W$ for all i . Then,

$$\begin{aligned} P(S_j|D) &= \sum_i P(S_j|E_i \wedge D)P(E_i|D) \\ &= \sum_i P(S_j|E_i \wedge D)q_i \text{ (by condition (1))} \\ &= \sum_i P(S_j|E_i)q_i \text{ (by condition (2))}, \end{aligned}$$

as required. \square

In fact, the theorem says that Bayesian conditioning in the enlarged algebra of propositions is in agreement with updating by MRE on a whole partition of propositions that belongs to some subalgebra of the enlarged one. This agreement rests on two conditions, originally introduced in Skyrms [29]. Condition (1) is an application of condition M, whereas Condition (2) is a kind of probabilistic independence called by Skyrms sufficiency. Both conditions have an intuitive appealing. Condition (1) says that the agent's prior degree of belief in E_i , conditional on the proposition specifying posterior degrees of belief over the members of \mathcal{E} , should be equal to the posterior degree of belief in E_i . This condition can be understood as saying that learning described by D is legitimate or justified. For example, it indicates that such a learning is not a result of memory loss. Sufficiency tells us that S_j is conditionally independent of D given each member of \mathcal{E} . Intuitively, if the agent knows which member of \mathcal{E} is true, then her knowledge about degrees of belief assigned to each member of that partition should have no bearing on her degree of belief in A . However, we should not regard these conditions as universally correct. Clearly, Condition (1) does not hold in epistemically "pathological" situations. Just consider the example of Ulysses and the sirens. Before hearing the siren's song, Ulysses has a high degree of belief that sailing among the rocks is dangerous. However, he is also sure that after hearing the sirens, he would cease to believe (wrongly as he now thinks) that sailing among the rocks is dangerous. If he were to obey Condition (1), he would have to cease to believe now that sailing among the rocks is dangerous. However, he now believes that this is not so, and so, Condition (1) is violated. Likewise, sufficiency does not hold in situations where S_j is a proposition X_{q_i} . Then, since D implies X_{q_i} , it is E_i , not D , that is irrelevant to S_j . However, whenever these two conditions hold, which seems to be fairly common, Bayesian conditioning in an enlarged degrees-of-belief space yields the same result as the MRE shift over a whole partition in the original smaller degrees-of-belief space.

Where does this result leave us *vis-à-vis* the question of whether a MRE shift on a whole partition satisfies the value of learning? To address this question, we first need to face a potential difficulty. Recall that in Good's argument, the experiment is represented by a finite partition of propositions, whose members are measurable subsets in W . However, the outputs of learning experiences represented by Constraint (i) are the values of posterior degrees of belief, not propositions. If this is so, how could the MRE updater assign degrees of belief to them? Furthermore, how could she determine the values of informed and uninformed decisions? By virtue of the representation introduced above, this difficulty can be mitigated by acknowledging that such values of posterior degrees of belief can be expressible as proposition D , which is a measurable subset in W . That is, from the point of view of the enlarged degrees-of-belief space, what we learn from the experiment reported by Constraint (i) is a proposition about the values of posterior degrees of belief over the members of a partition. Now, by moving to an enlarged degrees-of-belief space, we can think of a cost-free experiment as r possible results prompting r possible redistribution of the agent's degrees of belief over \mathcal{E} . Denote the m -th redistribution of the kind by the proposition D_m .

Now, it is easy to observe that by virtue of the representation captured in Theorem 1, the MRE updater on \mathcal{E} satisfies condition M, and thus, the value of learning theorem can be established. Since she can be represented as a Bayesian conditionalizer in the enlarged degree-of-belief space, in which the D_m 's are measurable subsets, her prior degree of belief for each state of the world S will be the expectation of her posterior degrees of belief in each S . These posteriors are given by the conditional prior degrees of

belief, the $P(A|D_m)$'s, defined in the enlarged degrees-of-belief space. More precisely, a demonstration that such an MRE update satisfies the value of learning may proceed as follows. The present value of making an uninformed decision is:

$$\begin{aligned} \max_i \sum_j P(S_j)U(A_i \wedge S_j) &= \max_i \sum_m \sum_j P(S_j|D_m)P(D_m)U(A_i \wedge S_j) \\ &= \max_i \sum_m \sum_j \frac{P(D_m|S_j)P(S_j)}{P(D_m)}P(D_m)U(A_i \wedge S_j) \\ &= \max_i \sum_m \sum_j P(S_j)P(D_m|S_j)U(A_i \wedge S_j). \end{aligned} \tag{10}$$

The posterior value of making a decision informed by D_m is given by:

$$\max_i \sum_j P(S_j|D_m)U(A_i \wedge S_j). \tag{11}$$

Given Equation (11), the present value of making a decision conditional on D_m is calculated by:

$$\begin{aligned} &\sum_m P(D_m) \max_i \sum_j P(S_j|D_m)U(A_i \wedge S_j) \\ &= \sum_m P(D_m) \max_i \sum_j \frac{P(D_m|S_j)P(S_j)}{P(D_m)} \\ &= \sum_m \max_i \sum_j P(S_j)P(D_m|S_j)U(A_i \wedge S_j), \end{aligned} \tag{12}$$

which is the prior expectation of the posterior value of making an informed decision. Now, it is easy to notice that on the same mathematical grounds as in Good's argument, the value given by Equation (12) is at least as great as the value given by Equation (10). Hence, MRE updating on \mathcal{E} represented as BC* is expected to be helpful and never harmful to one's decisions.

5. When the Value of Learning May Not Hold for MRE Updating

In this section, we examine the question of whether MRE updating in response to Constraint (ii) leads to the value of learning theorem. As will be apparent, the answer to this question is: it depends on how broadly one's learning experience reported by Constraint (ii) is described. More specifically, we show that whether the value of learning can be established in this case may be dependent on whether or not the contextual information, not reported by Constraint (ii), is taken into account in addition to the explicit information. We illustrate this point by means of the Judy Benjamin problem. If only the explicit information is taken into account in this case, then the value of learning theorem may not hold. By taking the contextual information into account, the constraint is made "complete" in the sense explicated below, and the value of learning theorem holds.

In general, consider a learning experience in which the agent learns the following conditional information "If A , then the odds for B are $\sigma/(1 - \sigma) : 1$ ", for $\sigma \in [0, 1]$. This information may prompt a change in the agent's conditional prior degrees of belief. That is, after learning this conditional information, her conditional prior degree of belief, $P(B|A)$, changes to her conditional posterior degree

of belief, $P'(B|A)$, which should be set equal to σ . With this constraint, we associate a closed and convex set of posterior degrees-of-belief functions $\mathcal{P}_\sigma = \{P' : P'(B|A) = \sigma\}$. In order to answer the question of whether a shift from P to P' , which belongs to this set and minimizes RE, leads to the value of learning theorem, we examine whether such a shift satisfies condition M.

For concreteness, we focus on the famous Judy Benjamin case, originally introduced in [5]. In this case, private Judy Benjamin is dropped in an area that is divided into two territories, the red territory (R) and the blue territory ($\neg R$). Each of these territories is further divided into the second company area (S) and headquarters company area ($\neg S$). These divisions form four quadrants. Initially, Judy assigns to each of the four quadrants a degree of belief of one quarter: $P(R \wedge S) = P(R \wedge \neg S) = P(\neg R \wedge S) = P(\neg R \wedge \neg S) = \frac{1}{4}$. Judy, then, receives the following radio message: “I don’t know where you are. If you are in the red territory, the odds are 3:1 that you are in the headquarters company area”. That is, the radio message prompts a change in one of Judy’s conditional degrees of belief by setting $P'(\neg S|R) = \frac{3}{4}$. Suppose further that Judy is an MRE updater, and let $R \wedge S, R \wedge \neg S, \neg R \wedge S, \neg R \wedge \neg S$ be the minimal elements of \mathcal{F} . Now, we may distinguish two ways of describing the constraint on Judy’s posterior degrees-of-belief function:

- (i*) The constraint pertains to all propositions of the partition $\{R \wedge S, R \wedge \neg S, \neg R\}$ of the elements in \mathcal{F} ;
- (ii*) The constraint pertains to some propositions of the partition $\{R \wedge S, R \wedge \neg S, \neg R\}$ of the elements in \mathcal{F} .

Let us consider each of these in turn. Case (i*) rests on the assumption that the MRE updater can obtain additional information about her posterior degrees of belief over the members of the entire partition by looking at the context of the Judy Benjamin case. The only explicit information she gets is the information about her posterior conditional degree of belief in $\neg S$, given R , *i.e.*, $P'(\neg S|R) = \frac{3}{4}$. Given this information, she knows how to set her posterior conditional degree of belief in S , given R : since all of her conditional degrees of belief must sum to one, we have that $P'(S|R) = \frac{1}{4}$. However, this does not yet provide a redistribution over the entire partition. What about her shift from $P(\neg R)$ to $P'(\neg R)$? This information is not given explicitly. However, this information can be gleaned from the context of the case: since the radio message does not say whether Judy is in the red or in the blue territory, it follows that her degree of belief in $\neg R$ remains unchanged, *i.e.*, $P'(\neg R) = P(\neg R) = P(\neg R \wedge S) + P(\neg R \wedge \neg S) = \frac{1}{2}$. This completes her redistribution over the entire partition of propositions. Let us assume that Judy’s learning experience does not lead her to the revision of her degree of belief in R . We thereby assume a condition called by Bradley [7] independence. Then, the sum of Judy’s posterior degrees of belief in $R \wedge \neg S$ and $R \wedge S$ equals her prior degree of belief in R , *i.e.*, $P'(R \wedge \neg S) + P'(R \wedge S) = P'(\neg S|R)P(R) + P'(S|R)P(R)$. Now, Judy’s task is to find the posterior degrees-of-belief function $P' \in \{P' : P'(R \wedge \neg S) = \frac{3}{8}, P'(R \wedge S) = \frac{1}{8}, P'(\neg R) = \frac{1}{2}\}$ that minimizes RE relative to P . As shown in [8] in a more general setting, RE is minimized iff for all $A \in \mathcal{F}$:

- (1) $P'(A|R \wedge \neg S) = P(A|R \wedge \neg S)$,
- (2) $P'(A|R \wedge S) = P(A|R \wedge S)$,
- (3) $P'(A|\neg R) = P(A|\neg R)$.

That is, Judy’s new degrees-of-belief function minimizes RE relative to her prior degree-of-belief function iff the shift in her degrees of belief is rigid and, thus, goes in accord with Jeffrey’s rule on the partition $\{R \wedge S, R \wedge \neg S, \neg R\}$. As emphasized in [8], by using the contextual information in the Judy Benjamin case, we can complete the constraint reporting Judy’s experience in a way that allows us to redistribute her new degrees of belief over the entire partition of propositions and to apply Jeffrey’s rule. Where does this result leave us *vis-à-vis* the question of whether an MRE shift in response to Constraint (ii) leads to the value of learning theorem? If Constraint (ii) pertains to the entire partition of propositions to which Jeffrey’s rule can be applied, then, in view of the representation given in Section 4, the MRE updater may be represented as a Bayesian conditionalizer in a degrees-of-belief space in which this constraint is a measurable subset of W . Consequently, she satisfies condition M, and thus, the value of learning holds for this case.

Things change if we turn to Case (ii*). Here, the radio message received by Judy prompts an incomplete redistribution of her degrees of belief over $\{R \wedge S, R \wedge \neg S, \neg R\}$. Here, we assume that no information, that makes the redistribution complete, can be gleaned from the context of this case. The radio message is the sole constraint imposed on her posterior degree-of-belief function. This explicit constraint causes her redistribution over $R \wedge S$ and $R \wedge \neg S$, leaving her posterior degree of belief in $\neg R$ unknown. However, as shown in [5], by using MRE updating, we can determine Judy’s posterior degree of belief in this proposition. However, this determination leaves us with a highly counter-intuitive consequence: $P'(\neg R) > \frac{1}{2}$, and hence, $P'(\neg R) > P(\neg R)$. That is, Judy’s new degree of belief in $\neg R$ is greater than her prior degree of belief $\neg R$, even if the radio message yields no information relevant to whether she is in the red rather than in the blue territory. More generally, for any value of σ , one’s posterior degree of belief in $\neg R$ that minimizes RE would be greater than one’s prior degree of belief in $\neg R$, and it remains unchanged only if $\sigma = \frac{1}{2}$. However, apart from being counter-intuitive, this observation shows that the MRE updater cannot satisfy condition M. To show this, we explore a result, due to Seidenfeld [3] and rehearsed by Uffink [30], which shows that MRE updating cannot be represented as Bayesian conditioning in an enlarged space in which an incomplete Constraint (ii) is a measurable subset of W , unless the constraint is irrelevant to one’s prior degrees of belief in $\neg R$. Suppose that Γ_σ (in the Judy Benjamin case, $\sigma = \frac{3}{4}$) is a measurable subset of W . Since, for any value of σ , the posterior degree of belief in $\neg R$ increases unless $P(\neg R) = P'(\neg R)$, we have that in the enlarged degrees-of-belief space:

$$P(\neg R) \geq \int_0^1 P(\neg R|\Gamma_\sigma)P(\Gamma_\sigma) d\sigma, \tag{13}$$

with strict inequality when there is some probability mass function on Γ_σ for $\sigma \neq 1/2$. That is, the prior degree of belief in $\neg R$ cannot be a convex combination of the conditional degrees of belief, the $P(\neg R|\Gamma_\sigma)$ ’s, for $\sigma \neq 1/2$. Not only does it show that MRE updating in Case (ii*) cannot be represented as Bayesian conditioning in the enlarged space, but also it shows that MRE updating in that case fails to satisfy condition M, unless $P(\neg R) = P'(\neg R)$. If the conditional degrees of belief, the $P(\neg R|\Gamma_\sigma)$ ’s, are understood as possible posterior degrees of belief, the $P'(\neg R)$ ’s, then we have that:

$$P(\neg R) \geq \sum_{P'} P'(\neg R)P(X_{P'}) \tag{14}$$

Consequently, $P(\neg R|X_{P'}) \geq P'(\neg R)$, and so, condition M does not hold in general. Additionally, given that condition M is both necessary and sufficient for the value of learning to hold, it follows that MRE updating does not in general lead to the value of learning theorem. That is, MRE updating may lead to a decrease in expected utility.

The above analysis has an interesting philosophical import. Whether MRE updating leads to the value of learning theorem in the case of Constraint (ii) crucially depends on whether or not the agent takes into account the contextual information. However, this should not strike us as odd; for there is nothing in the machinery of MRE updating that could determine the unique way of describing one's learning experience. This opens the possibility of using both explicit and contextual information in order to determine a given constraint. More to the point, MRE does not suffice to guarantee the value of learning when the new information comes as constraints over conditional degrees of belief. It has been shown that to guarantee the value of learning, MRE must be supplemented by some additional rule, which tells us how to add extra constraints gleaned from the context.

Note, however, that Case (ii*) also points towards another notion of context sensitivity. This has to do with how MRE determines the lacking information about one's posterior degree of belief in $\neg R$. Though this information is not given explicitly, MRE could fill in the blanks for us. However, whether it does this adequately depends on the details of a given learning situation, which also include the context. On the widespread view, in the Judy Benjamin case, MRE does not fill in the blanks adequately, for it leads to counter-intuitive results: after updating, Judy's degree of belief in $\neg R$ increases, while it should remain unchanged. However, it is perfectly possible to add to the Judy Benjamin case a story indicating that the choice of the blue or red territory is dependent on the choice of the red headquarters company area or the red second company area. However, this type of context-sensitivity should be distinguished from the one described above. For whatever story we plot in the Judy Benjamin case, MRE may provide us with the lacking information in a way that violates condition M, as indicated by [13]. In contrast, the type of context sensitivity we alluded to above has consequences for whether or not condition M is satisfied by the MRE updater.

Let us point out some consequences of our analysis. The fact that in some cases the application of MRE and its justification in terms of the value of learning is context-sensitive lends credence to the idea that updating rules are essentially tools in the "art of judgment" rather than universally valid inductive rules. In this spirit, Bradley [7] (p. 362) points out that even Bayes's rule "should not be thought of as a universal and mechanical rule of updating, but as a technique to be applied in the right circumstances, as a tool in what Jeffrey terms the 'art of judgment'". Similarly, Douven and Romeijn [8] (p. 660) stress that adopting an updating rule based on minimizing distance between degrees of belief to cover updating on conditional information "may be an art, or a skill, rather than a matter of calculation or derivation from more fundamental epistemic principles". Our analysis shows that even a justification of MRE updating in terms of the value of learning cannot proceed mechanically. Rather, it requires a careful consideration of the entire learning experience that the agent undergoes.

It is important to emphasize that our analysis should not be regarded as providing a support to yet another idea, widely discussed within the degrees-of-belief dynamics, called by van Fraassen [31] voluntarism. According to this idea, deliverances of experience should be understood as commands that constrain the agent's posterior degrees of belief. These commands reflect the agent's decision to accept

whatever her learning experience reveals. It is not hard to observe that voluntarism may lead to the idea that belief change is sensitive to what the agent accepts as her constraint. After all, two agents may accept different constraints on their posterior degrees of belief, even if they undergo the same learning experience. However, this is different from saying that the way in which we respond to a constraint depends on the context of our learning experience; for the context is not a feature of the agent's epistemic attitudes, but rather, it is a part of the learning experience that bears on the agent's epistemic attitudes. Hence, whether or not the context of a given case contributes to one's learning experience is not a matter of one's voluntary decision. Of course, according to voluntarism, the agent might voluntarily decide not to take the contextual information as her constraint. However, our analysis does not force us to accept this possibility.

6. Concluding Remarks

Clearly, our analysis is not a full story on the justification of MRE in terms of the value of learning. We have discussed this issue with respect to only two types of constraints: the first pertaining to a redistribution of one's degrees of belief over the entire partition of propositions; the second pertaining to a change in one's conditional degrees of belief. Despite this limitation, we have shown that the justification of MRE updating is not so simple a task as one might think. By fitting MRE updating and Bayesian conditioning together in an enlarged space, we have shown that in cases involving the first constraint, MRE leads to the value of learning. However, we have argued that this might not be so in cases involving the second type of constraint. In such cases, whether or not the value of learning holds crucially depends on whether the context of one's learning experience is taken into account.

We may transfer the insights of our analysis to the discussion about the status of MRE updating. Recall that initially, we have distinguished, from various views on this issue, the view on which MRE updating is universally valid and the views that deny its universal validity. It is tempting to think that if this rule of updating were universally valid, it would be neutral with respect to how a given learning experience is described. Moreover, it seems that if it were universally valid, its justification would not depend on whether or not the contextual information is reported by a given constraint. The findings of this paper show that neither the application of MRE nor its justification are so neutral. Hence, they lend credence to the claim that MRE is not a universal or mechanical updating rule.

Acknowledgments

I would like to thank Jan-Willem Romeijn for stimulating conversations on the subject matter of this paper and for his incredibly helpful comments on earlier versions of this paper. Also, I want to thank the editor and the two referees for valuable comments on this paper. This research was supported by a Vidi grant from the Netherlands Organization of Scientific Research (NWO grant 016114354: "What are the chances?").

Conflicts of Interest

The author declares no conflict of interest.

References

1. Williams, P.M. Bayesian conditionalisation and the principle of minimum information. *Br. J. Philos. Sci.* **1980**, *31*, 131–144.
2. Friedman, K.; Shimony, A. Jaynes' maximum entropy prescription and probability theory. *J. Stat. Phys.* **1971**, *4*, 381–384.
3. Seidenfeld, T. Entropy and uncertainty. *Philos. Sci.* **1986**, *53*, 467–491.
4. Shimony, A. The status of the principle of maximum entropy. *Synthese* **1985**, *63*, 35–53.
5. Van Fraassen, B.C. A problem for relative information minimizers in probability kinematics. *Br. J. Philos. Sci.* **1981**, *32*, 375–379.
6. Van Fraassen, B.C.; Hughes, R.I.G.; Harman, G. A problem for relative information minimizers, continued. *Br. J. Philos. Sci.* **1986**, *37*, 453–463.
7. Bradley, R. Radical probabilism and Bayesian conditioning. *Philos. Sci.* **2005**, *72*, 342–364.
8. Douven, I.; Romeijn, J.-W. A new resolution of the Judy Benjamin problem. *Mind* **2011**, *120*, 637–670.
9. Skyrms, B. Maximum entropy inference as a special case of conditionalization. *Synthese* **1985**, *63*, 55–74.
10. Skyrms, B. Updating, supposing, and maxent. *Theory Dec.* **1987**, *22*, 225–246.
11. Shore, J.; Johnson, R. Properties of cross-entropy minimization. *IEEE Trans. Inf. Theory* **1981**, *27*, 472–482.
12. Shore, J.E.; Johnson, R.W. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Trans. Inf. Theory* **1980**, *26*, 26–37.
13. Grünwald, P. Maximum entropy and the glasses you are looking through. In Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI 2000), Stanford University, Stanford, CA, USA, 30 June–3 July 2000; Morgan Kaufmann Publishers: Burlington, MA, USA, 2000; pp. 238–246.
14. Leitgeb, H.; Pettigrew, R. An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philos. Sci.* **2010**, *77*, 236–272.
15. Savage, L.J. *The Foundations of Statistics*; Jon Wiley and Sons, Inc.: New York, NY, USA, 1954.
16. Good, I.J. On the principle of total evidence. *Br. J. Philos. Sci.* **1967**, *17*, 319–321.
17. Huttegger, S.M. Learning experiences and the value of knowledge. *Philos. Stud.* **2014**, *171*, 279–288.
18. Graves, P.R. The total evidence theorem for probability kinematics. *Philos. Sci.* **1989**, *56*, 317–324.
19. Skyrms, B. *The Dynamics of Rational Deliberation*; Harvard University Press: Cambridge, MA, USA, 1990.
20. Skyrms, B. The structure of radical probabilism. *Erkenntnis* **1997**, *45*, 285–297.
21. Gaifman, H. A Theory of Higher Order Probabilities. In *Causation, Chance, and Credence*; Skyrms, B., Harper, W., Eds.; Morgan Kaufmann Publishers: San Francisco, CA, USA, 1988; pp. 191–219.
22. Van Fraassen, B.C. Belief and the will. *J. Philos.* **1984**, *81*, 235–256.

23. Weisberg, J. Conditionalization, reflection, and self-knowledge. *Philos. Stud.* **2007**, *135*, 179–197.
24. Easwaran, K. Expected accuracy supports conditionalization—and conglomerability and reflection. *Philos. Sci.* **2013**, *80*, 119–142.
25. Diaconis, P.; Zabell, S.L. Updating subjective probability. *J. Am. Stat. Assoc.* **1982**, *77*, 822–830.
26. Jeffrey, R. *The Logic of Decision*, 2nd ed.; University of Chicago Press: Chicago, IL, USA, 1983.
27. Grünwald, P.D.; Halpern, J.Y. Updating probabilities. *J. Artif. Intell. Res.* **2003**, *19*, 243–278.
28. Skyrms, B. *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*; Yale University Press: New Haven, CT, USA, 1980.
29. Skyrms, B. Higher order degrees of belief. In *Prospects for Pragmatism*; Mellor, D.H., Ed.; Cambridge University Press: Cambridge, UK, 1980; pp. 109–137.
30. Uffink, J. The constraint rule of the maximum entropy principle. *Stud. Hist. Philos. Sci. B* **1996**, *27*, 47–79.
31. Van Fraassen, B.C. *Laws and Symmetry*; Clarendon Press: Oxford, UK, 1989.

© 2015 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/4.0/>).