Generalized Learning and Conditional Expectation

Simon M. Huttegger   Michael Nielsen

Abstract

Reflection and martingale principles are central to models of rational learning. They can be justified in a variety of ways. In what follows we study martingale and reflection principles in the context of measure theory. We give special attention to two approaches for justifying these principles that have not been studied in that context before: diachronic coherence and the value of information. Together with an extant argument based on expected accuracy, these arguments lend support to the thesis that reflection and martingale principles govern rational learning.

1 Introduction

Richard Jeffrey's epistemological program known as *radical probabilism* (Jeffrey, 1992) tries to reconcile two goals: it aims at being both mathematically precise and flexible regarding representations of opinions and opinion change. *Jeffrey conditioning* (also known as *probability kinematics*) is a prime example (Jeffrey, 1957, 1965, 1968). Being less rigid than standard Bayesian conditioning, it nonetheless captures a common type of uncertain observational learning within the framework of probability theory.

Learning from experience is not, however, restricted to Bayesian and Jeffrey conditioning. In the most extreme case the learning event is a black box: we only know an agent's opinions before and after the learning experience. Despite the lack of structure, there is a family of principles—known as *reflection* and *martingale principles*—that are plausible candidates for regulating how to respond to new information in generalized learning situations. The principles say, roughly speaking, that current opinions should cohere with anticipated future opinions.

How can such principles be justified? There are three paradigms (Huttegger, 2017). One is based on diachronic Dutch book arguments (Goldstein, 1983; van Fraassen, 1984; Skyrms, 1990); the second one proceeds in terms of expected accuracy (Easwaran, 2013; Huttegger, 2013); and the third approach appeals to ideas related to the value of knowledge theorem in decision theory (Huttegger, 2014). The three paradigms are well understood in finite probability spaces. The same is true for the expected accuracy approach and infinite probability spaces. However, a corresponding treatment is missing for the other two approaches. The aim of our paper is to fill this gap.

We are going to follow the most common approach to infinite probability spaces, Kolmogorov's measure-theoretic probability theory (Kolmogorov, 1933, 1956). The central concept of Kolmogorov's theory is the *conditional expectation of a random variable given a σ-algebra*, which gives rise to a natural formulation of generalized learning. There is an ongoing discussion of the limits of the concept of conditional expectation and the measure-theoretic approach. Kolmogorov's assumption of countable additivity has not found universal acceptance (de Finetti, 1974; Seidenfeld, 2001). Moreover, conditional expectation has some strange consequences when it is taken with respect to certain σ-algebras (Seidenfeld et al., 2001). For the purposes of this paper, we take it for granted that conditional expectation is a worthwhile object of study, even if it may not be universally valid.

The paper is structured as follows. After introducing Kolmogorov conditional expectation in §2, we introduce reflection and martingale principles in §3. Most of the paper investigates the value-of-knowledge-paradigm (§4). In §5 we show that the diachronic Dutch book argument for Kolmogorov conditional probability presented in Rescorla (2018) gives rise reflection and martingale principles. Finally, in §6 we conclude by discussing the impact of these arguments on the status of those principles.

2 Conditional Probability and Conditional Expectation

Let $(\Omega, \mathfrak{B}, \mathbb{P})$ be a probability space: Ω is a set of atomic events, $\mathfrak{B}$ is a σ-algebra of measurable subsets of Ω called "events," and $\mathbb{P}$ is a countably additive probability measure.[1] We can think of $(\Omega, \mathfrak{B}, \mathbb{P})$ as an experimental setup, where the elements of Ω represent the most fine-grained experimental outcomes, members of $\mathfrak{B}$ capture experimental events, and $\mathbb{P}$ represents a scientist's degrees of belief.

Let $\mathfrak{G}$ be a sub-σ-algebra of $\mathfrak{B}$. Intuitively, $\mathfrak{G}$ represents the outcomes of an experiment that gives partial information about the space of atomic events. Furthermore, let $X : \Omega \to \mathbb{R}$ be a random variable on $(\Omega, \mathfrak{B})$. Random variables represent quantities that have a determinate value in the experimental setup (i.e., X is $\mathfrak{B}$-measurable).

According to Kolmogorov (1933), the *conditional expectation* of X given $\mathfrak{G}$, $\mathbb{E}[X|\mathfrak{G}]$, is itself a random variable. More precisely, $\mathbb{E}[X|\mathfrak{G}]$ is a $\mathfrak{G}$-measurable real-valued function s.t.

$$\int_G X d\mathbb{P} = \int_G \mathbb{E}[X \mid \mathfrak{G}] d\mathbb{P}$$

[1] Here we set aside issues that arise when $\mathbb{P}$ is finitely but not countably additive. See Kadane et al. (1996) and Seidenfeld (2001).

for all $G \in \mathfrak{G}$. The conditional expectation $\mathbb{E}[X|\mathfrak{G}]$ is almost surely unique and can be thought of as the best estimate of X given the experimental outcomes in $\mathfrak{G}$.[2] For instance, if $\mathfrak{G}$ is the trivial algebra $\{\emptyset, \Omega\}$, then $\mathfrak{G}$ provides no information at all, and the best estimate of X given $\mathfrak{G}$ is the expectation of X, $\mathbb{E}[X|\mathfrak{G}] = \mathbb{E}[X]$, a.s. (almost surely). On the other hand, $\mathbb{E}[X|\mathfrak{B}] = X$ a.s., since X is $\mathfrak{B}$-measurable and thus coincides with its best estimate given $\mathfrak{B}$.

It is well known that Kolmogorov's concept of conditional expectation basically coincides with the traditional concept of conditional expectation when Ω is finite. Moreover, *conditional probability* can be defined in terms of conditional expectation. Let B be an event in $\mathfrak{B}$, and let χ_B be the indicator of B. Then the conditional probability of B given $\mathfrak{G}$, $\mathbb{P}[B|\mathfrak{G}]$, is equal to $\mathbb{E}[\chi_B|\mathfrak{G}]$ a.s. While conditional probability, thus understood, is a random variable, it effectively coincides with the elementary concept of conditional probability when Ω is finite.

3 Principles for Updating Opinions

Conditioning is the best known Bayesian principle of updating. It is usually stated in terms of partitions: the new probability of an event A is given by the conditional probability of A given the element of the partition that was observed to be true. Conditional probabilities given a σ-field generalize that idea to epistemic situations which cannot be captured by partitions. Conditional expectations given a σ-field generalize Bayesian conditioning to random variables.

In both manifestations, Bayesian conditioning assumes that one learns something with certainty.[3] Learning with certainty is a limiting case of *uncertain learning*, as Jeffrey has argued in several places (Jeffrey, 1957, 1965, 1968). This led him to develop a broader Bayesian epistemology, *radical probabilism*, which subsumes Jeffrey conditioning, but also includes other forms of updating (Jeffrey, 1992).

Several authors have argued that there is a family of principles, ususally referred to as reflection or martingale principles, that govern how an agent should change opinions in response to new information in all forms of updating that fall under the umbrella of radical probabilism (Goldstein, 1983; van Fraassen, 1984; Skyrms, 1990; Jeffrey, 1992). Our first order of business is to formulate them within the conceptual framework sketched in the preceding section.[4]

[2] See Huttegger (2013) for a brief discussion. The treatment presupposes strictly proper scoring rules. For brevity's sake we cannot go into details here.

[3] This understanding requires conditional probabilities given a σ-field to be proper. See our discussion of properness below.

[4] These principles have been met with criticism (e.g. Levi, 1987; Talbott, 1991; Maher, 1992; Bovens, 1995; Briggs, 2009). We refer the reader to Huttegger (2017, Chapters 5 and 6) for an extensive discussion

As before, $(\Omega, \mathfrak{B}, \mathbb{P})$ is a probability space. Let $P : \Omega \times \mathfrak{B} \to [0, 1]$ be a function with the following properties:

(i) $P(\omega, \cdot)$ is a probability measure on $(\Omega, \mathfrak{B})$ for each $\omega \in \Omega$;

(ii) $P(\cdot, B)$ is $\mathfrak{B}$-measurable for each $B \in \mathfrak{B}$.

We will refer to P as an *update policy*. An update policy represents an agent's future probabilities after learning. Let $\mathfrak{F}$ be the sigma-algebra generated by the family $\{P(\cdot, B) : B \in \mathfrak{B}\}$. Note that $\mathfrak{F}$ is a sub-σ-algebra of $\mathfrak{B}$, since by (ii) $P(\cdot, B)$ is $\mathfrak{B}$-measurable. Intuitively, $\mathfrak{F}$ represents information about future probabilities after learning. It is the "future algebra."

The *principle of reflection* says that present probabilities, conditional on the future, should be equal to future probabilities, or, for all $B \in \mathfrak{B}$:

$$\mathbb{P}(B \mid \mathfrak{F}) = P(\cdot, B) \text{ a.s.} \tag{1}$$

In other words, for all $B \in \mathfrak{B}$, $P(\cdot, B)$ is a version of the conditional probability $\mathbb{P}(B \mid \mathfrak{F})$. Notice that reflection *implies* that $\mathbb{P}(\cdot \mid \mathfrak{F})$ has a regular version. The latter is true, for example, if Ω is a Polish space and $\mathfrak{B}$ is its Borel σ-algebra. In general, however, Ω's being Polish is not necessary for the existence of a regular version of $\mathbb{P}(\cdot \mid \mathfrak{F})$. Conversely, if a space has no regular conditional probability given $\mathfrak{F}$, then the principle of reflection does not hold.

Suppose the reflection principle holds. By integrating both sides of (1), we get

$$\mathbb{P}(B) = \int P(B) d\mathbb{P}. \tag{2}$$

We call (2) the *weak reflection principle*, since it is implied by the reflection principle.

The martingale principle generalizes the reflection principle to $\mathbb{P}$-integrable random variables. In the remainder of the paper, we denote the space of $\mathbb{P}$-integrable random variables by $\mathcal{L}^1$. The *martingale principle* says that, for all $X \in \mathcal{L}^1$ for which $\int X dP \in \mathcal{L}^1$:

$$\mathbb{E}[X|\mathfrak{F}] = \int X dP \text{ a.s.} \tag{3}$$

If $X = \chi_B$ is an indicator function, then $\int X dP = P(B)$ is bounded and therefore integrable, so the martingale principle implies the reflection principle. In general, it is not the case for an arbitrary random variable in $X \in \mathcal{L}^1$ that $\int X dP \in \mathcal{L}^1$ as well. We will

of these issues based on the idea that the reflection principle and the martingale principle apply to learning events and not arbitrary kinds of opinion change.

soon show, however, that if the weak reflection principle holds, then $\int X dP$ is indeed in $\mathcal{L}^1$.

We interpret the martingale principle as saying that *best estimates* after learning are (almost surely) equal to conditional expectations given the future algebra. This interpretation requires viewing $\int X dP$ as the best estimate for X after learning. This view seems quite intuitive to us. Since coherent best estimates for random variables are always given by expected values, the view amounts to assuming that learning maintains coherence. To spell this out in more detail, it seems very plausible that an agent's best estimate for X should be the same as her fair price for a gamble with X. It is the price she would pay (respectively, accept) to have the opportunity to collect (respectively, pay out) whatever value the random variable X manifests in the actual world. If the agent's best estimates did not match her fair prices, then she would be making a mistake by her own lights. She would regard her own betting behavior as falling short of the standard set by her *best* estimates. With best estimates and fair prices thus identified, suppose now that the agent expects to remain coherent after learning. (It isn't clear in what sense a change from coherent beliefs to incoherent ones could count as learning, but we can set this issue aside.) By the standard Dutch book argument, coherence demands that fair prices be determined by expected values (e.g., Jeffrey, 2004, Chapter 4). So, since best estimates match fair prices, the agent's best estimate for X after learning is given by the expected value of X relative to her update policy P. This argument suggests that, if an agent expects to remain coherent after learning, then her best estimate for X after learning should be given by the P-expected value of X. If, moreover, the martingale principle holds, then best estimates after learning are given by conditional expectations.[5]

The martingale principle implies

$$\int X d\mathbb{P} = \int \int X dP d\mathbb{P} \tag{4}$$

for all $X \in \mathcal{L}^1$ for which $\int X dP \in \mathcal{L}^1$. This says that random variables and their best estimates after learning have the same expected value. We refer to (4) as the *weak martingale principle.*

Our first result verifies that $\int X dP$ is integrable in the presence of weak reflection and also shows that the (weak) martingale principle is actually equivalent to the (weak) reflection principle.

[5]These arguments are pragmatic. One might wonder if there is also an accuracy approach to establishing that best estimates must be expectations. For squared error loss this was worked out for previsions by de Finetti (1974). There are other loss functions one might plausibly use in certain contexts. For the sake of brevity we won't say more about this here.

Proposition 1. *If the weak reflection principle holds, then $\int XdP$ is $\mathfrak{F}$-measurable and integrable for all $X \in \mathcal{L}^1$. Moreover, the reflection principle is equivalent to the martingale principle, and the weak reflection principle is equivalent to the weak martingale principle.*

Proof. First note that by considering $X = \chi_B$, it is clear that the (weak) martingale principle implies the (weak) reflection principle, so we only need to verify the converse implications to settle the proposition's second claim.

Now, suppose that the weak reflection principle holds. Establishing the proposition's first claim and deriving the weak martingale principle is a straightforward application of what Williams (1991) calls "the standard machine" : we first show that the desired results holds if X is an indicator function, then if X is a simple function, then if X is non-negative, and finally if X is any integrable random variable.

If $X = \chi_B$ is an indicator function, then clearly $\int XdP = P(B)$ is $\mathfrak{F}$-measurable, and (4) follows immediately from (2). If $X = \sum_{i=1}^n b_i \chi_{B_i}$ is a simple function, then $\int XdP$ is a linear combination of $\mathfrak{F}$-measurable functions, and so is itself $\mathfrak{F}$-measurable; moreover, $\int XdP$ is clearly bounded and therefore integrable, and

$$\int Xd\mathbb{P} = \sum_{i=1}^n b_i \mathbb{P}(B_i) = \sum_{i=1}^n b_i \int P(B_i) d\mathbb{P} = \int \int XdPd\mathbb{P},$$

by the linearity of the integral and the weak reflection principle. If X is non-negative, then it can be approximated from below by an increasing sequence of simple functions $Z_1, Z_2, \ldots$. Then, by the monotone convergence theorem, $\int XdP$ is the monotone limit of a sequence of $\mathfrak{F}$-measurable functions, and so itself $\mathfrak{F}$-measurable, and

$$\int Xd\mathbb{P} = \lim_{n\to\infty} \int Z_n d\mathbb{P} = \lim_{n\to\infty} \int \int Z_n dPd\mathbb{P} = \int \int XdPd\mathbb{P},$$

by the monotone convergence theorem and our previous step. This last equation also reveals that $\int XdP$ is integrable, because X is. Finally, for arbitrary $X \in \mathcal{L}^1$, write $X = X^+ - X^-$.[6] Since, X^+ and X^- are non-negative and integrable, by our last step we have that both $\int X^+ dP$ and $\int X^- dP$ are integrable. It follows that $\int XdP = \int X^+ dP - \int X^- dP$, being the difference of two $\mathfrak{F}$-measurable integrable functions, is itself $\mathfrak{F}$-measurable and integrable. Moreover,

$$\begin{aligned}\int Xd\mathbb{P} &= \int X^+ d\mathbb{P} - \int X^- d\mathbb{P} = \int \int X^+ dPd\mathbb{P} - \int \int X^- dPd\mathbb{P} \\ &= \int \Big(\int X^+ dP - \int X^- dP \Big) d\mathbb{P} = \int \int XdPd\mathbb{P}.\end{aligned}$$

[6]Recall that $X^+ = \max(X, 0)$ and $X^- = \max(-X, 0)$.

We have established the proposition's first claim and shown that the weak reflection principle implies the weak martingale principle. The argument showing that the reflection principle implies the martingale principle is analogous and we omit it. □

4 Value of Information

In a seminal paper, I. J. Good (1967) showed that in certain finite decision problems a Bayesian agent will prefer making a decision after conditioning on new information to making a decision right away. This result, known as the *value of knowledge theorem*, was anticipated by F. P. Ramsey and L. J. Savage (Ramsey, 1990; Savage, 1954). Good used it as an argument for the *principle of total evidence*, which says that in forming opinions one should use all available evidence. Skyrms (1990) generalized Good's theorem to Polish spaces. His argument extends to all probability spaces that admit a regular conditional probability.

The value of knowledge theorem operates under a number of stringent assumptions. One is that there are no costs involved in obtaining new evidence. In addition, the agent is assumed to conform to the principles of Savage's decision theory before and after the learning event, and the learning event has no effect on the utilities associated with decision problems. There are a number of additional assumptions that we won't mention here (see Skyrms, 1990, for details).

We are not so much interested in defending the plausibility of the assumptions underlying the value of knowledge theorem, but in using the value of knowledge idea to study Bayesian learning from experience. For this we tacitly assume that the assumptions hold. Huttegger (2014) showed that the value of information can be used to derive the principle of reflection when Ω is finite. Huttegger's central requirement is that a belief change leads one to expect to make better decisions in a sufficiently large class of decision problems if the belief change is a *genuine learning event*, in which the agent responds to legitimate new information.

Let's articulate this idea in the context of a probability space $(\Omega, \mathfrak{B}, \mathbb{P})$ and an update policy P. An *act* is an integrable real-valued random variable $A : \Omega \to \mathbb{R}$. The value $A(\omega)$ represents the utility of state ω if act A is chosen. Utilities are assumed to be cardinal utilities. The expected utility of A is thus given by the expectation of A:

$$\int A(\omega)\mathbb{P}(d\omega).$$

A *decision problem* $\mathcal{A}$ is a finite set of acts.

Prior to a learning event, the expected value of choosing an act from a decision problem

$\mathcal{A} = \{A_i\}$ is:

$$\max_i \int A_i d\mathbb{P}$$

We refer to this as the *value of the prior Bayes act*, since it is the utility a Bayesian agent expects to get in case she makes a choice before new information is available.

After the learning event, the agent adopts the probabilities given by her update policy P. Provided that $\int A_i dP$ is integrable, the *prior expected value of the posterior Bayes act is* given by the following integral:

$$\int \max_i \int A_i(\omega)P(\omega', d\omega)\mathbb{P}(d\omega').$$

We now suppose that the update policy P is in response to a genuine learning event. Then the *value of information postulate* says that for all decision problems $\mathcal{A} = \{A_i\}$, $\int A_i dP$ is integrable and

$$\int \max_i \int A_i dP d\mathbb{P} \geq \max_i \int A_i d\mathbb{P}. \tag{5}$$

This says that the agent does not expect to make worse decisions after acquiring new information.[7]

The next result shows that the value of information postulate (5) is equivalent to the weak reflection principle (2).

Proposition 2. *The value of information postulate is equivalent to both the weak reflection principle and the weak martingale principle.*

Proof. By Proposition 1, it suffices to establish the equivalence with the weak reflection principle. If the value of information postulate holds, i.e.

$$\int \max_i \int A_i dP d\mathbb{P} \geq \max_i \int A_i d\mathbb{P}$$

for all decision problems $\mathcal{A} = \{A_i\}$, then it holds in particular for $\mathcal{A} = \{\chi_B\}$ with $B \in \mathfrak{B}$. Hence,

$$\int P(B) d\mathbb{P} \geq \mathbb{P}(B) \tag{6}$$

holds for all $B \in \mathfrak{B}$. If the inequality were strict for some B, then putting B^c in (6) would yield a contradiction, so (6) holds with equality for all B, which is to say that the weak reflection principle (2) holds.

[7] The postulate in Huttegger (2014) is actually a bit stronger since it includes a proviso to the effect that in a certain class of cases the weak inequality in (5) is strict.

Now suppose that the weak reflection principle holds. By Proposition 1, the weak martingale principle holds as well, and $\int A_i dP \in \mathcal{L}^1$ for all decision problems $\mathcal{A} = \{A_i\}$. Let $A_{\max}$ denote a maximizer of $\int A_i d\mathbb{P}$. Then,

$$\max_i \int A_i dP \geq \int A_{\max} dP$$

Integrating both sides of this, we obtain

$$\int \max_i \int A_i dP d\mathbb{P} \geq \int \int A_{\max} dP d\mathbb{P} = \int A_{\max} d\mathbb{P} = \max_i \int A_i d\mathbb{P},$$

by the weak martingale principle.

□

We will now discuss two ways of improving Proposition 2. We provide two conditions under which the value of information postulate is equivalent to the reflection principle.

Let us say that P is *proper* just in case $P(\omega, F) = 1_F(\omega)$ for all ω in a $\mathbb{P}$-probability 1 set and all $F \in \mathfrak{F}$. Elementary conditional probabilities are trivially proper in the sense that $\mathbb{P}(E \mid E) = 1$ whenever $\mathbb{P}(E) > 0$. In general, however, regular conditional probabilities are not proper.[8] Properness corresponds to the intuition that conditioning on $\mathfrak{F}$ represents observing which members of $\mathfrak{F}$ are true and which are false. The condition that P is proper is thus one way to generalize the "luminosity" condition that Huttegger (2014) used for finite probability spaces, which requires that the agent perfectly recalls how learning affected her probabilities.

We regard the next proposition as one of the two main results of this paper. We shall return to a discussion of its significance in the final section.

Proposition 3. *If P is proper, then the value of information postulate holds if and only if both the reflection principle and the martingale principle hold.*

Proof. By Proposition 1, it suffices to establish the equivalence with the reflection principle. Suppose that P is proper and that the value of information postulate holds. By Proposition 2, this implies that the weak reflection principle (2) holds. Moreover, almost surely, for all $F \in \mathfrak{F}$, if $\omega \in F$, $P(\omega, F) = 1$, which implies $P(\omega, B \cap F) = P(\omega, B)$ for all $B \in \mathfrak{B}$. That is, $\chi_F P(B \cap F) = \chi_F P(B)$ a.s. Similarly, if $\omega \notin F$, then $P(\omega, F) = 0$, which implies $P(\omega, B \cap F) = 0$. So, $\chi_{F^c} P(B \cap F) = 0$ a.s. Then, by the weak reflection principle, for all $B \in \mathfrak{B}$ and all $F \in \mathfrak{F}$,

$$\mathbb{P}(B \cap F) = \int P(B \cap F) d\mathbb{P} = \int_F P(B \cap F) d\mathbb{P} + \int_{F^c} P(B \cap F) d\mathbb{P} = \int_F P(B) d\mathbb{P}.$$

[8] See Blackwell and Dubins (1975) and Seidenfeld et al. (2001).

This shows that for all $B \in \mathfrak{B}$, $P(\cdot, B)$ is a version of $\mathbb{P}(B \mid \mathfrak{F})$; that is, the reflection principle holds.

Conversely, whether or not P is proper, the reflection principle implies the weak reflection principle, which in turn implies the value of information postulate, by Proposition 2. □

The connection between the principle of reflection and the value of information postulate is further elucidated by the following property. We say that P is *idempotent* if for $\mathbb{P}$ almost every ω and all $B \in \mathfrak{B}$:

$$P(\omega, B) = \int P(\omega', B) P(\omega, d\omega'). \tag{7}$$

The meaning of idempotence becomes clearer if the integral is interpreted as representing the probability of an event. Then the right-hand side of (7) is the probability of first updating states in Ω and then updating the probability of B. Idempotence requires that the resulting probability is the same as updating the probability of B immediately. Thus, iterating P does not change the result of updating beyond the initial application.[9]

Regular conditional probabilities are idempotent. So, if the reflection principle holds, then so do weak reflection and idempotence. We establish the converse with the next proposition.

Proposition 4. *If P is idempotent and the weak reflection principle holds, then both the reflection principle and the martingale principle hold.*

Proof. We establish that the martingale principle holds. In the same way that one can extend the (weak) reflection principle to the (weak) martingale principle, as in the proof of Proposition 1, so we can extend idempotence to

$$\int X(\omega') P(\omega, d\omega') = \int \int X(\omega'') P(\omega', d\omega'') P(\omega, d\omega') \tag{8}$$

for all $X \in \mathcal{L}^1$.

Now, we will show that there exists a sub-σ-algebra $\mathfrak{G}$ of $\mathfrak{B}$ such that, for all $X \in \mathcal{L}^1$, $\int X dP$ is a version of $\mathbb{E}[X|\mathfrak{G}]$. Let L^1 be the quotient space of $\mathbb{P}$-integrable random variables that are equal with $\mathbb{P}$-probability 1. By Corollary 1 in Douglas (1965), as presented by Gyenis and Redei (2017, Proposition 2.7), it suffices to show that the mapping $\mathbb{Y} : L^1 \to L^1, X \mapsto \int X dP$ is a linear projection, contractive in the L^1 norm, and such that $\mathbb{Y}(\chi_\Omega) = \chi_\Omega$. All of these properties are basically immediate from the definitions of the terms involved.

[9] Idempotence was first studied in the context of Markov chains by Blackwell (1942).

By Proposition 1, $\mathbb{Y}(X) \in L^1$ for all $X \in L^1$. The linearity of the integral implies that $\mathbb{Y}$ is linear. It is clear that $\mathbb{Y}(\chi_\Omega) = \chi_\Omega$. To say that $\mathbb{Y}$ is a projection means that $\mathbb{Y}(\mathbb{Y}(X)) = \mathbb{Y}(X)$, and this follows from the definition of $\mathbb{Y}$ and (8). To say that $\mathbb{Y}$ is contractive in the L^1 norm means that $\mathbb{E}|\mathbb{Y}(X)| \leq \mathbb{E}|X|$, and this follows from the fact that $|\mathbb{Y}(X)| \leq \mathbb{Y}(|X|)$ and (4), which holds by Proposition 1.

Thus,

$$\mathbb{E}[X|\mathfrak{G}] = \mathbb{Y}(X) \tag{9}$$

for some sub-σ-algebra $\mathfrak{G}$ and all $X \in L^1$. Thus, for all $X \in \mathcal{L}^1$,

$$\mathbb{E}[X|\mathfrak{G}] = \int X dP \text{ a.s.} \tag{10}$$

We conclude by showing that (10) continues to hold with $\mathfrak{G} = \mathfrak{F}$. Since $\int XdP$ is measurable in both $\mathfrak{G}$ and $\mathfrak{F}$, it follows from (10) that $\int XdP$ is a version of $\mathbb{E}[X|\mathfrak{G} \cap \mathfrak{F}]$:

$$\mathbb{E}[X \mid \mathfrak{G} \cap \mathfrak{F}] = \mathbb{E}[\mathbb{E}[X \mid \mathfrak{G}] \mid \mathfrak{G} \cap \mathfrak{F}] = \mathbb{E}\Big[\int XdP \mid \mathfrak{G} \cap \mathfrak{F}\Big] = \int XdP \text{ a.s.}$$

By the minimality of $\mathfrak{F}$, $\mathfrak{F} \subseteq \mathfrak{G} \cap \mathfrak{F}$. Hence, $\mathfrak{F} = \mathfrak{G} \cap \mathfrak{F}$, and we are done. □

The next proposition is the second main result of the paper.

Proposition 5. *Idempotence and the value of information postulate hold if and only if both the reflection principle and martingale principle hold.*

Proof. Proposition 4 shows that idempotence and weak reflection imply the reflection principle. We already observed the converse above. Since, by Proposition 2, the value of information principle is equivalent to the weak reflection principle, the result follows. □

Let $C_\omega = \{\omega' : P(\omega', B) = P(\omega, B) \text{ for all } B \in \mathfrak{B}\}$. C_ω is the event that updated probabilities are given by $P(\omega, \cdot)$. Then, idempotence is connected to the following *luminosity condition*:

$$P(\omega, C_\omega) = 1 \text{ for a.e. } \omega. \tag{11}$$

This says that the update policy does not assign positive probability to the event that it updates in a different way than it had planned.[10] The update policy, in other words, doesn't forget any information that was gained from learning.

In order for (11) to be meaningful, C_ω must be measurable. To guarantee measurability, we assume that $\mathfrak{B}$ is *countably generated*. In that case, $\mathfrak{F}$ is countably generated, and

[10] This is the analog to the luminosity condition that was used in Huttegger (2014).

$\{C_\omega : \omega \in \Omega\}$ is a family of $\mathfrak{F}$ atoms in $\mathfrak{F}$.[11] To see why, let $\mathfrak{A}$ be a countable algebra generating $\mathfrak{B}$. Then $\mathfrak{F} = \sigma(\{P(\cdot, B) : B \in \mathfrak{A}\})$ by a monotone class argument.[12] Thus, $\mathfrak{F}$ is countably generated: for each $B \in \mathfrak{A}$, $\sigma(P(\cdot, B))$ is countably generated (call the generating set $\mathcal{C}_B$) because $P(\cdot, B)$ is a real-valued random variable, and $\mathfrak{F}$ is generated by $\cup_{B\in\mathfrak{A}}\mathcal{C}_B$, which is a countable union of countable sets. Moreover, it follows that $C_\omega = \bigcap_{B\in\mathfrak{A}}\{\omega' : P(\omega, B) = P(\omega', B)\}$ because $\mathfrak{A}$ is an algebra generating $\mathfrak{B}$, and, since this intersection is countable, $C_\omega \in \mathfrak{F}$. The collection of all C_ω may or may not be countable. In fact, one can see that there are countably many C_ω just in case the range of P—i.e., the collection of probability measures $\{P(\omega, \cdot) : \omega \in \Omega\}$—is countable.

The following results relate luminosity to the conditions studied above. The first follows from a theorem of Blackwell (1942) and was stated independently by Gaifman (1988).

Proposition 6. *Suppose $\mathfrak{B}$ is countably generated. Then the luminosity condition implies idempotence. If the weak reflection principle holds, then the luminosity condition and idempotence are equivalent.*

Proof. Suppose the luminosity condition holds. Then, almost surely, $P(\omega, C_\omega) = 1$, and thus

$$\int P(\omega', B)P(\omega, d\omega') = \int P(\omega', B)1_{C_\omega}(\omega')P(\omega, d\omega').$$

Now, for all $\omega' \in C_\omega$, $P(\omega, \cdot) = P(\omega', \cdot)$, and so the integrand on the right side is equal to $P(\omega, B)$. The conclusion follows.

The second statement follows from Blackwell (1942, Theorem 7), which implies that there is a set N such that $P(\omega, N) = 0$ for (almost) all ω and such that (11) holds for (almost) all $\omega \notin N$. Weak reflection implies that $\mathbb{P}(N) = 0$. □

By Proposition 5 and Proposition 6 we have the following.

Corollary 1. *Suppose $\mathfrak{B}$ is countably generated. The luminosity condition and the value of information postulate hold if and only if both the reflection principle and martingale principle hold.*

It remains to clarify the relationship between the luminosity condition and proper update policies.

[11] An atom in $\mathfrak{F}$ is a set of the form $\bigcap_{\omega\in A\in\mathcal{A}} A$, ω fixed. If C_ω is an atom in $\mathfrak{F}$, then for every $F \in \mathfrak{F}$ either $C_\omega \cap F = C_\omega$ or $C_\omega \cap F = \emptyset$ (Blackwell, 1942, p. 563).

[12] Let $\mathfrak{G} = \sigma(\{P(\cdot, B) : B \in \mathfrak{A}\})$ and note that the class of all measurable sets B such that $P(\cdot, B)$ is $\mathfrak{G}$-measurable is a monotone class since measurability is preserved under monotone limits. By the monotone class theorem, this class is identical to $\sigma(\mathfrak{A}) = \mathfrak{B}$. Thus, $P(\cdot, B)$ is $\mathfrak{G}$-measurable for every $B \in \mathfrak{B}$.

Proposition 7. *If $\mathfrak{B}$ is countably generated, then properness and luminosity are equivalent.*

Proof. Suppose $\mathfrak{B}$ is countably generated. Then the C_ω are the atoms of $\mathfrak{F}$, and for every $F \in \mathfrak{F}$ either $C_\omega \cap F = C_\omega$ or $C_\omega \cap F = \emptyset$. Suppose the luminosity condition holds, and let $F \in \mathfrak{F}$. Then, for all ω in a set of $\mathbb{P}$-probability one, if $\omega \in F$, then $\omega \in C_\omega \subset F$ and $P(\omega, F) = 1$; if $\omega \notin F$, then $C_\omega \cap F = \emptyset$, so $P(\omega, F) = 0$. Therefore properness holds. Conversely, since $\mathfrak{B}$ is countably generated, $C_\omega \in \mathfrak{F}$; and since $\omega \in C_\omega$, properness implies that $P(\omega, C_\omega) = 1$ for all ω in a set of $\mathbb{P}$-probability one. □

5 Diachronic Dutch Books

We now turn to reflection and martingale principles in the context of dynamic coherence. Rescorla (2018) has shown that P is a conditional probability given a sub-σ-algebra $\mathfrak{G}$ if and only if P is not vulnerable to a diachronic Dutch book when events in $\mathfrak{G}$ are observed. The σ-algebra $\mathfrak{G}$ captures a learning situation in which the agent updates on the outcome of an experiment or an observational setup. This is only superficially different from the learning situations considered in this paper. Rescorla's argument extends straightforwardly to our setting.

For the reflection principle, replace Rescorla's observational σ-algebra $\mathfrak{G}$ with our future algebra $\mathfrak{F}$. His argument shows that $\mathbb{P}[\cdot|\mathfrak{F}]$ is the almost surely unique updating strategy that avoids a diachronic Dutch book. More precisely, Rescorla's main results imply that there is no diachronic Dutch book for the update policy P if and only if $P(\cdot, B)$ is a version of $\mathbb{P}[B|\mathfrak{F}]$ for every $B \in \mathfrak{B}$; that is, if and only if the reflection principle (1) holds and, by Proposition 1 if and only if the martingale principle holds.

The Dutch book characterization of reflection, then, simply generalizes the well-known results of Goldstein (1983) and van Fraassen (1984) to measure-theoretic conditional probabilities and expectations: in a black-box learning situation, diachronic coherence demands conditioning on future probabilities. This point is also discussed in Skyrms (1997, p. 287).

6 Concluding Remarks

The two main results of this paper can be summarized as follows:

(i) In the presence of properness, the value of knowledge postulate is equivalent to the principle of reflection and the martingale principle (Proposition 3).

(ii) The value of knowledge postulate and idempotence are equivalent to the principle of reflection and the martingale principle (Proposition 5).

One way to read these results is as justification of the reflection and martingale principle. The value of knowledge postulate is a rather compelling requirement for all genuine learning situations: in the absence of factors that interfere adversely with the acquisition of new information, one should not expect to make worse decisions after adjusting one's probabilities. Properness is a desirable property for update policies. An improper update policy foresees that it will lose some information that is part of the learning event. While information loss is certainly something that can happen after a learning experience, it does not seem to be entirely appropriate for update policies. If the update policy foresees an information loss, then this should already be factored into the learning event in such a way that the future algebra only captures the actual information gain.

Still, luminosity conditions such as properness have sometimes been criticized as unreasonable assumptions about our introspective access to belief change (Weisberg, 2007; Williamson, 2002). Idempotence establishes an alternative route to the principle of reflection. It simply requires that updating be independent of whether information is presented in one or more steps. Unlike properness, idempotence holds for all regular conditional probabilities; it thus constitutes a plausible constraint for all update policies.

Our results can also be read as putting in context so-called counterexamples to the principle of reflection. Any violation of reflection entails a violation of the value of knowledge postulate or a violation of properness and idempotence. This identifies the conceptual space within which we can find counterexamples. They draw on situations in which an agent does not expect to make better decisions based on new beliefs, in which some information that was learned is lost, or in which belief change depends on the way in which information is presented.

Finally, many criticisms of the principle of reflection go hand in hand with criticisms of dynamic Dutch book arguments. Based on our results and Huttegger (2017), critics of reflection face a more formidable task, though. Skepticism regarding the principle of reflection does not just entail skepticism regarding diachronic coherence but also accuracy and value of information arguments. Such skepticism might prove difficult to sustain in the face of the general treatment of those arguments given here and, for expected accuracy, in Huttegger (2013).

References

Blackwell, David. 1942. "Idempotent Markoff Chains." *Annals of Mathematics* 43:560–567.

Blackwell, David, and Lester E. Dubins. 1975. "On Existence and Non-Existence of Proper, Regular, Conditional Distributions." *The Annals of Probability* 3:741–752.

Bovens, Luc. 1995. "'P and I Will Believe That Not-P': Diachronic Constraints on Rational Belief." *Mind* 104:737–760.

Briggs, R. A. 2009. "Distorted Reflection." *Philosophical Review* 118:59–85.

de Finetti, Bruno. 1974. *Theory of Probability*, Volume 1. London: John Wiley & Sons.

Douglas, Ronald G. 1965. "Contractive Projections on an $\mathfrak{L}_1$ Space." *Pacific Journal of Mathematics* 15:443–462.

Easwaran, Kenny. 2013. "Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection." *Philosophy of Science* 80:119–142.

Gaifman, Haim. 1988. A Theory of Higher Order Probabilities. In *Causation, Chance and Credence*, ed. Brian Skyrms and William L. Harper, 191–219. Dordrecht: Kluwer.

Goldstein, Michael. 1983. "The Prevision of a Prevision." *Journal of the American Statistical Association* 78:817–819.

Good, I. J. 1967. "On the Principle of Total Evidence." *British Journal for the Philosophy of Science*, 17:319–321.

Gyenis, Zalán and Redei, Miklos. 2017. "General Properties of Bayesian Learning as Statistical Inference Determined by Conditional Expectations." *The Review of Symbolic Logic* 10:719–755.

Huttegger, Simon M. 2013. "In Defense of Reflection." *Philosophy of Science* 80:413–433.

— 2014. "Learning Experiences and the Value of Knowledge." *Philosophical Studies* 171:279–288.

— 2017. *The Probabilistic Foundations of Rational Learning.* Cambridge: Cambridge University Press.

Jeffrey, Richard C. 1957. "Contributions to the Theory of Inductive Probability." PhD diss., Princeton University.

— 1965. *The Logic of Decision.* New York: McGraw-Hill, New York. 3rd revised edition Chicago: University of Chicago Press, 1983.

— 1968. "Probable Knowledge." In *The Problem of Inductive Logic*, ed. Imre Lakatos, 166–180. Amsterdam: North-Holland.

— 1992. *Probability and the Art of Judgement.* Cambridge: Cambridge University Press.

— 2004. *Subjective Probability. The Real Thing*. Cambridge: Cambridge University Press.

Kadane, Jay B., Mark J. Schervish, and Teddy Seidenfeld. "Reasoning to a Foregone Conclusion." *Journal of the American Statistical Association* 91:1228–1235.

Kolmogorov, A N. (1933). *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer, Berlin.

Kolmogorov, Andrey N. 1956. *Foundations of the Theory of Probability*, volume 25. New York: Chelsea Publishing Company.

Kreps, David M. 1988. *Notes on the Theory of Choice*. Boulder: Westview Press.

Levi, Isaac. 1987. "The Demons of Decision." *The Monist* 70:193–211.

Maher, Patrick. 1992. "Diachronic Rationality." *Philosophy of Science* 59:120–141.

Ramsey, Frank P. 1990. "Weight or the Value of Knowledge." *British Journal for the Philosophy of Science* 41:1–4.

Rescorla, Michael 2018. "A Dutch Book Theorem and Converse Dutch Book Theorem for Kolmogorov conditionalization. *The Review of Symbolic Logic* 11:705–735.

Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: Dover Publications.

Seidenfeld, Teddy. 2001. "Remarks on the Theory of Conditional Probability: Some Issues of Finite Versus Countable Additivity." In *Probability Theory*, ed. Vincent F. Hendricks, 167–178. Dordrecht: Kluwer.

Seidenfeld, Tedyy, Mark J. Schervish, and Jay B. Kadane. 2001. "Improper Regular Conditional Distributions." *Annals of Probability* 29:1612–1624.

Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.

Skyrms, Brian. 1997. "The Structure of Radical Probabilism." *Erkenntnis* 45:285–297.

Talbott, William J. 1991. "Two Principles of Bayesian Epistemology." *Philosophical Studies* 62:135–150.

van Fraassen, Bas C. 1984. "Belief and the Will." *Journal of Philosophy* 81:235–256.

Weisberg, Jonathan. 1991. "Conditionalization, Reflection, and Self-Knowledge." *Philosophical Studies* 135:179–197.

Williams, David. 1991. *Probability With Martingales.* Cambridge: Cambridge University Press.

Williamson, Timothy. 2002. *Knowledge and its Limits.* Oxford: Oxford University Press.