The Review of Symbolic Logic
Volume 14, Number 4, December 2021

A NEW ARGUMENT FOR KOLOMOGOROV CONDITIONALIZATION

MICHAEL NIELSEN
Department of Philosophy, University of Sydney

Abstract. This paper contributes to a recent research program that extends arguments supporting elementary conditionalization to arguments supporting conditionalization with general, measure-theoretic conditional probabilities. I begin by suggesting an amendment to the framework that Rescorla (2018) has used to characterize regular conditional probabilities in terms of avoiding Dutch book. If we wish to model learning scenarios in which an agent gains complete membership knowledge about some subcollection of the events of interest to her, then we should focus on updating policies that are what I shall call *proper*. I go on to characterize regular conditional probabilities in proper learning scenarios using what van Fraassen (1999) calls The General Reflection Principle.

§1. Introduction. What is the correct way to revise one's probability judgments in the light of newly acquired information? It depends, of course, on the kind of information that one acquires. Suppose an agent acquires information that allows her to determine the true member of a finite set of exhaustive and mutually exclusive propositions, each of which she currently assigns positive probability. Then, the widely held Bayesian view is that she should update her probability judgments by what I shall call *elementary conditionalization*. Elementary conditionalization says that the agent's revised probability judgments should be given by her initial conditional probability judgments, given the proposition that she learns is true. If $\mathbb{P}$ represents the agent's initial probability judgments, P her updated probability judgments, and E is the proposition that she learns is true, then elementary conditionalization says that the updated probabilities should be given by $P = \mathbb{P}(\cdot \mid E)$, where the right-hand side is defined whenever $\mathbb{P}(E) > 0$.

Why use elementary conditionalization? There are several arguments supporting it. I will mention three. First, it can be shown that elementary conditionalization is the only updating policy that ensures invulnerability to diachronic Dutch books in certain learning scenarios (Teller, 1973). Second, agents who update their probability judgments by elementary conditionalization maximize the expected accuracy of those judgments (Leitgeb & Pettigrew, 2010a, 2010b; Greaves & Wallace, 2006). Third, decision makers who update by elementary conditionalization increase their expected utility gains (Good, 1967).

Received: March 2, 2019.
2010 *Mathematics Subject Classification*: 60A05, 60A10.
Key words and phrases: conditionalization, Kolmogorov, regular conditional probability, reflection principle.

© Association for Symbolic Logic, 2020
doi:10.1017/S1755020320000179

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

But despite the arguments supporting elementary conditionalization, it is clearly a limited updating policy. For information need not, and often does not, come in the form of a finite partition whose members are assigned positive initial probability. Any learning scenario in which an agent observes the value of a continuous random variable, for instance, cannot be handled by elementary conditionalization. Such learning scenarios are ubiquitous in scientific practice. A Bayesian account of these sorts of scenarios requires tools that are more sophisticated and more general than elementary conditional probabilities. One such tool was defined by Kolmogorov (1950) and is nowadays called a *regular conditional probability*. Regular conditional probabilities have become a standard part of mathematical probability theory, but, with a few exceptions, have received little attention from philosophers.

In response to this neglect, there is an ongoing effort in philosophy to generalize the arguments supporting elementary conditionalization to arguments supporting conditionalization with regular conditional probabilities. Rescorla (2018) has proved Dutch book theorems that characterize this more general kind of conditionalization. Huttegger (2013) has shown that accuracy-based arguments can be generalized as well. Huttegger & Nielsen (2019) extend the decision-theoretic arguments of Good (1967).[1]

The present paper contributes to this research program by generalizing an argument for elementary conditionalization due to van Fraassen (1999). The argument shows that, in a certain class of learning scenarios, conditionalization is the only correct updating policy. I will begin, in §2, by clarifying what a learning scenario is. Then I will define the class of learning scenarios that delimits the paper's argument. I call these learning scenarios *proper*. The definition of properness is obtained by amending Rescorla's (2018) framework. I will motivate this amendment by arguing that it captures the type of learning scenario in which conditionalization is usually taken to apply. In §3, I define *Kolmogorov conditionalization*, which generalizes elementary conditionalization. In §4, I present and discuss general versions of van Fraassen's reflection principle and another principle similar to it that I call Mixing. Reflection and Mixing are then used to characterize Kolmogorov conditionalization in proper learning scenarios (§5). After proving the characterization result, I will discuss how it gives rise to an argument for Kolmogorov conditionalization. §6 contains some additional mathematical observations, and §7 concludes.

§2. Learning scenarios. A learning scenario is a situation in which an agent expects to acquire new information that may prompt her to change her beliefs, probability judgments, preferences, and other attitudes. I will be focused exclusively on learning's impact on probability judgments, setting aside questions about how to revise attitudes like preferences. The aim of this section is to define a particular class of learning scenarios, what Rescorla (2018) calls *Kolmogorov learning scenarios*. I will begin by motivating Kolmogorov learning scenarios informally, and then I will define them formally. My exposition in this part of the section follows Rescorla quite closely. Later in the section, however, I will argue that we ought to distinguish between two kinds of Kolmogorov learning scenario in order to properly model the informal considerations that motivate us at the outset. To be clear, this point is not intended as a criticism of

[1] Also see Gyenis, Hofer-Szabo, & Rédei (2017) and Gyenis & Rédei (2017) for accounts of Bayesian learning that use general conditional probabilities and conditional expectations.

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

Rescorla's framework, but rather a friendly amendment to it. I return to this point at the end of the section.

A Kolmogorov learning scenario is a formal model of the following situation. An agent begins an inquiry with some initial probability judgments about a collection of events that she is interested in. She expects her inquiry—or some particular stage of it—to deliver *partial information* about some subcollection of these events. In particular, she expects to find out which events in the subcollection are true and which are false. The agent has an *updating policy*, which is a plan for revising her initial probability judgments upon receiving the partial information. In other words, her update policy is a recipe for transforming her initial probability judgments given knowledge about which members of the subcollection are true and which are false. The question before us is, What kind of updating policy should an agent in this scenario adopt?

Rescorla (2018) has shown that an agent with certain pragmatic concerns must adopt an updating policy given by a conditional probability. If the agent wants to guarantee that she will not adopt an updating policy that will lead her to accept a collection of bets that ensures her a loss of utility, then she must adopt an updating policy that instructs her to transform her initial probability judgments by conditioning them on the partial information that she receives.

I will now introduce Kolmogorov learning scenarios, which model the informal remarks above. Since this model of learning has been discussed at length in a number of recent publications, my treatment here will be brief.[2]

An agent's initial probability judgments are modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Ω is a set of elementary outcomes and $\mathcal{F}$ is a sigma-algebra of subsets of Ω, the members of which are called *events*. $\mathbb{P}$ is a countably additive probability measure. Until §6, all probability measures are assumed to be countably additive. The probability space represents the agent's uncertainty about which $\omega \in \Omega$ is the true outcome. Partial information about Ω is modeled by a sigma-subalgebra $\mathcal{G}$ of $\mathcal{F}$.

We are interested in learning scenarios in which the agent gains membership knowledge about $\mathcal{G}$. As Rescorla puts it, "In Kolmogorov's setup, the agent learns whether ω [the true outcome] belongs to each $G \in \mathcal{G}$" (2018, p. 709). Put in other words, the agent learns which events in $\mathcal{G}$ are true and which are false. Rescorla suggests that this feature of the learning scenario is modeled by a property of the agent's updating policy, to which we now turn.

Recall that, intuitively, an updating policy is the agent's plan for revising her initial probability judgments in light of the partial information given by $\mathcal{G}$. Formally, an *updating policy* is a function $P : \mathcal{F} \times \Omega \to [0, 1]$ of two variables that has the following properties:

(i) For each $\omega \in \Omega$, $P(\cdot, \omega)$ is a probability measure on $(\Omega, \mathcal{F})$;
(ii) For each $A \in \mathcal{F}$, $P(A, \cdot)$ is a $\mathcal{G}$-measurable function.

For each $A \in \mathcal{F}$ and $\omega \in \Omega$, think of $P(A, \omega)$ as the agent's updated probability for A given the partial information about the outcome ω that is contained in $\mathcal{G}$. Or, in Rescorla's words, $P(A, \omega)$ is the probability for A given ω "as filtered through $\mathcal{G}$" (2018, p. 10).

To make the definition of an updating policy clearer, let me say a bit more about the properties (i) and (ii). The first simply says that the agent's updating policy

[2] In addition to Rescorla's (2018) detailed development, see, for instance, Huttegger (2015).

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

transforms initial probability judgments into *probability* judgments. We are setting aside updating policies that instruct agents to adopt incoherent judgments. Property (ii) is less transparent. In general, to say that a function $f : \Omega \to [0, 1]$ is $\mathcal{G}$-measurable means that, for all $r \in [0, 1]$, the set of points that f maps to a number no greater than r, i.e. $\{\omega \in \Omega : f(\omega) \leq r\}$, is a member of $\mathcal{G}$. What this means in the case of (ii) is that $\{\omega \in \Omega : P(A, \omega) \leq r\} \in \mathcal{G}$ for all events A and all $r \in [0, 1]$. And, because $\mathcal{G}$ is a sigma-algebra, the same holds upon reversing the inequality, replacing with strict inequalities, and replacing with equality. So, any question we might ask about the updating policy that is of the form, "In which outcomes will the revised probability of A be no greater than/strictly less than/equal to/etc. r?," will be answered by the partial information given by $\mathcal{G}$. To put this in yet another way, condition (ii) says that the partial information in $\mathcal{G}$ *determines* the updating policy's instructions for revising initial probability judgments. This is much like saying that revised probabilities are a *function* of the partial information learned, and nothing else.

A *Kolmogorov learning scenario*, then, is a quintuple $(\Omega, \mathcal{F}, \mathcal{G}, \mathbb{P}, P)$, where $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space, $\mathcal{G}$ is a sigma-subalgebra of $\mathcal{F}$, and P is an updating policy. Rescorla's main contribution is a pair of results that together characterize when P is a regular conditional probability of $\mathbb{P}$ given $\mathcal{G}$. Let us now recall—again, briefly—the definition of regular conditional probabilities.

A function $\mathbb{P}_{\mathcal{G}} : \mathcal{F} \times \Omega \to [0, 1]$ is called a *regular conditional probability* for $\mathbb{P}$ given $\mathcal{G}$ if and only if $\mathbb{P}_{\mathcal{G}}$ satisfies (i) and (ii)—which is to say, $\mathbb{P}_{\mathcal{G}}$ is an updating policy—as well as

(iii) For all $A \in \mathcal{F}$ and $G \in \mathcal{G}$,

$$\int_G \mathbb{P}_{\mathcal{G}}(A) d\mathbb{P} = \mathbb{P}(A \cap G),$$

which is a general form of the law of total probability. The term "regular" refers to the satisfaction of (i). For a fixed $A \in \mathcal{F}$, a $\mathcal{G}$-measurable function satisfying the integral equation in (iii) for all $G \in \mathcal{G}$ can be shown to exist. Such a function is called a *conditional probability* for A given $\mathcal{G}$, but we will not be focused on these objects.[3]

An updating policy given by elementary conditionalization is a special case of an update policy given by a regular conditional probability. Recall that elementary conditionalization applies to learning scenarios in which the agent expects to learn the true member of a positive partition of the outcome space. In the language of this section, the agent expects to gain membership knowledge about $\mathcal{G}$, where $\mathcal{G}$ is generated by a finite partition $\{E_1, \ldots, E_n\}$ such that $\mathbb{P}(E_i) > 0$ for all $i \in \{1, \ldots, n\}$. If the agent adopts an updating policy that satisfies (iii)—that is, if her updating policy is given by a regular conditional probability—then her updating policy is $P(A, \omega) = \mathbb{P}(\cdot \mid E(\omega))$, where $E(\omega)$ is whichever member of the partition $\{E_1, \ldots, E_n\}$ is true at ω. So, the model that we are developing captures elementary conditionalization as a special case.

Not all updating policies satisfy (iii), so not all updating policies are given by regular conditional probabilities. What Rescorla shows is that P is a regular conditional probability just in case it is invulnerable to certain kinds of diachronic Dutch book. The updating policies that deem unacceptable any collection of bets that leads to a sure loss are exactly the updating policies given by regular conditional probabilities.

[3] For more discussion, see Rescorla (2018, pp. 711–712, 728) and Seidenfeld (2001). We return to questions about the existence of regular conditional probabilities in the next section.

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

I will not engage with the details of Rescorla's Dutch book theorems here. Rather, I wish to return to an earlier remark. Above, following Rescorla, I motivated the formal definition of a Kolmogorov learning scenario by saying that, in such scenarios, agents gain membership knowledge about $\mathcal{G}$. At the outset, we were interested in learning scenarios in which the agent receives information that allows her to determine which members of $\mathcal{G}$ are true and which are false. If the agent learns which members of $\mathcal{G}$ are true and which are false, then whatever the true outcome ω is, the agent should be able to say, for each $G \in \mathcal{G}$, whether or not $\omega \in G$. In other words, if we are modeling a learning scenario in which the agent expects to gain membership knowledge about $\mathcal{G}$, then we are assuming that the agent expects learning to result in her being able to determine, for each $\omega \in \Omega$ and $G \in \mathcal{G}$, whether or not $\omega \in G$. But that is just to say that we are assuming the agent's updating policy has the following property:

(iv) For all $\omega \in \Omega$ and $G \in \mathcal{G}$, $P(G, \omega) = 1_G(\omega)$,

where 1_G is the indicator function for G, which returns 1 for $\omega \in G$ and 0 for $\omega \notin G$. If the agent adopts an updating policy that does not satisfy (iv), then this reflects a judgment that the learning scenario in which she finds herself will *not* deliver membership knowledge about $\mathcal{G}$. This could happen in a number of ways. For example, if the agent adopts an updating policy P such that $P(G, \omega) = 1/2$ for some ω and $G \in \mathcal{G}$, then this reflects a judgment that learning might not resolve uncertainty about every member of $\mathcal{G}$. If, to take another example, the agent adopts an updating policy such that $P(G, \omega) = 1$ for some $\omega \notin G$, then this reflects a judgment that learning might be *misleading* about some events in $\mathcal{G}$. In either case, failures of (iv) are incompatible with scenarios in which the agent expects to gain complete membership knowledge about $\mathcal{G}$.

The trouble now is, as we have defined them, Kolmogorov learning scenarios need not satisfy (iv). In particular, (iv) is not entailed by the conjunction of (i) and (ii). This means that our formal model does not really correspond to our motivating remarks. In fact, the model is more general than the motivating remarks suggest and captures learning scenarios in which the agent does *not* gain membership knowledge about $\mathcal{G}$. Moreover, (iv) is also not entailed by the conjunction of (i), (ii), and (iii). Regular conditional probabilities that satisfy (iv) are called *proper*, and it is well known that improper regular conditional probabilities exist.[4] This means that Rescorla's characterization of regular conditional probabilities in terms of avoiding Dutch book applies to a class of learning scenarios that is more general than the intended class. Rescorla's result holds even in some learning scenarios that do not provide membership knowledge about $\mathcal{G}$. To be more precise, the result continues to hold in scenarios where the agent's updating policy is given by an improper regular conditional probability.

Some of Rescorla's remarks suggest that he interprets (ii)—which says that an updating policy is $\mathcal{G}$-measurable in its second argument—as capturing the idea that learning delivers complete membership knowledge about $\mathcal{G}$: "Thus, $\mathcal{G}$-measurability captures the intuitive idea that someone who gains full membership knowledge for $\mathcal{G}$ regarding ω thereby gains implicit knowledge whether $\mathbb{P}_{\mathcal{G}}(A, \omega) \leq r$" (Rescorla, 2018, p. 711). But $\mathcal{G}$-measurability is a constraint that is quite distinct from the assumption that agents gain full membership knowledge about $\mathcal{G}$. As I hope my remarks above

[4] See Blackwell & Ryll-Nardzewski (1963), Blackwell & Dubins (1975), Dubins (1975), and Seidenfeld, Schervish, & Kadane (2001).

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

have indicated, $\mathcal{G}$-measurability corresponds to the idea that updating policies should be *determined* by $\mathcal{G}$; a policy for updating on the information in $\mathcal{G}$ that depends on information *not in* $\mathcal{G}$ is no policy for updating on $\mathcal{G}$ at all. In principle, though, an updating policy can be determined by the information in $\mathcal{G}$ without that information providing membership knowledge about $\mathcal{G}$.[5] For example, an agent might expect learning to deliver only probabilistic information about $\mathcal{G}$. These are the sorts of learning scenarios that probability kinematics was designed to handle (Jeffrey, 1965, 1992).[6] To make this more concrete, consider the following simple example. Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$, $G_1 = \{\omega_1, \omega_2\}$, and $G_2 = \{\omega_3\}$. Let $\mathcal{F}$ be the powerset of Ω and $\mathcal{G}$ be the sigma-subalgebra generated by the partition $\{G_1, G_2\}$. If $\omega \in G_1$, let $P(\cdot, \omega)$ be uniform on Ω, and let $P(\cdot, \omega_3)$ be point mass at ω_3. Then, $P(A, \cdot)$ is $\mathcal{G}$-measurable for every $A \in \mathcal{F}$, but it is not the updating policy of an agent who expects to gain complete membership knowledge about $\mathcal{G}$. At ω_1, for instance, there is still uncertainty about G_1 because $P(G_1, \omega_1) = 2/3$. Perhaps, however, Rescorla does not intend (ii) to correspond to the assumption that agents gain membership knowledge about $\mathcal{G}$. But if so, it is difficult to see what other aspect of the model corresponds to this assumption. The only other substantive constraint is (i), which only says that updating policies are coherent.

In view of the remarks above, I submit that it is worth exploring Kolmogorov learning scenarios that satisfy (iv). For they are the learning scenarios in which agents gain membership knowledge of $\mathcal{G}$, and therefore the scenarios that generalize the ones in which elementary conditionalization applies. Following the terminology introduced above, let us call such scenarios *proper* Kolmogorov learning scenarios, and let us say that an updating policy is proper when it satisfies (iv). My aim in the remainder of the paper is similar to Rescorla's, but my focus is narrower. I aim to characterize when updating policies are given by regular conditional probabilities in *proper* Kolmogorov learning scenarios.

In focussing on a narrower class of learning scenarios, I do not mean to suggest that updating policies *ought* to be proper. There may well be learning scenarios in which agents cannot reasonably expect to gain membership knowledge about a sigma-subalgebra of the outcome space, as I recently illustrated with a simple example. Nor do I mean to criticize Rescorla's more general framework. A characterization of regular conditional probabilities, whether proper or not, in terms of avoiding Dutch book is an important achievement. Dutch book theorems play a foundational role in Bayesian thinking (Ramsey, 1931; de Finetti, 1974), and it is crucial from a theoretical point of view to understand how the classical results extend to the measure-theoretic setting in which most contemporary probability theory takes place. The reason for focussing on proper learning scenarios is simply that they are particularly natural and interesting. They generalize the assumption underlying elementary conditionalization—that the agent gains membership knowledge of a finite partition—and, as van Fraassen (1989, 1999) has argued, they correspond to well-designed scientific experiments. Moreover, as we will see below, regular conditional probabilities admit a nice characterization in proper Kolmogorov learning scenarios in terms of what van Fraassen has called The General Reflection Principle.

[5] And, conversely, an agent can expect to gain membership knowledge about $\mathcal{G}$ without thereby adopting a $\mathcal{G}$-measurable updating policy. Although (iv) implies that $P(G, \cdot)$ is $\mathcal{G}$-measurable for all $G \in \mathcal{G}$, it does not imply (ii).

[6] Also see Huttegger (2015).

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

§3. Kolmogorov conditionalization. Before beginning to develop the main result of the paper, I will quickly give a more precise account of what it means to say that an updating policy P is given by a regular conditional probability. Recall from the previous section that a regular conditional probability $\mathbb{P}_{\mathcal{G}}$ is any function satisfying (i)–(iii). It's important to note now that such a function may not exist and, when such a function does exist, it may not be unique. This means that there are potentially many functions satisfying (i)–(iii) and potentially none at all. We shall deal with the existence problem by simply assuming that regular conditional probabilities exist. In fact, in most common probability spaces, regular conditional probabilities *do* exist, so this is a mild assumption.[7] The nonuniqueness of functions f satisfying (i)–(iii) is captured by saying that f is a *version* of $\mathbb{P}_{\mathcal{G}}$.[8] It should be noted that existence and uniqueness problems do not arise for elementary conditionalization because whenever $\mathcal{G}$ is generated by a finite, positive partition, $\mathbb{P}_{\mathcal{G}}$ exists and is unique.

The next definition clarifies the notion of Kolmogorov conditionalization that will be the focus of the rest of the paper.

Definition 1. *Let $(\Omega, \mathcal{F}, \mathcal{G}, \mathbb{P}, P)$ be a Kolmogorov learning scenario in which a regular conditional probability $\mathbb{P}_{\mathcal{G}}$ for $\mathbb{P}$ given $\mathcal{G}$ exists. Then, we say that P is given by* Kolmogorov conditionalization *if and only if P is a version of $\mathbb{P}_{\mathcal{G}}$.*

The main result of the paper is a generalization of results in van Fraassen (1999) and characterizes Kolmogorov conditionalization in proper Kolmogorov learning scenarios. I now start working towards that result by introducing the characterizing conditions.

§4. Mixing and the general reflection principle. In this section, a random variable $f : \Omega \to \mathbb{R}$ is a $\mathcal{F}$-measurable function. (The definition of $\mathcal{F}$-measurability is analogous to that of $\mathcal{G}$-measurability, discussed in §2.) If f is a bounded random variable on $(\Omega, \mathcal{F})$, let us say that the real number r is a *mixture of f* just in case $r = \int f d\mu$ for some probability measure μ on $(\Omega, \mathcal{F})$. Note that, if Ω is finite, r is a mixture of f in our sense just in case r lies in the interval spanned by $\{f(\omega) : \omega \in \Omega\}$; or, in other words, just in case r is a convex combination of $\{f(\omega) : \omega \in \Omega\}$. The next definition extends this terminology to updating policies.[9]

Definition 2. *Say that $\mathbb{P}$ is a* mixture *of the updating policy P and that P satisfies* Mixing *if and only if there exists a probability measure μ on $(\Omega, \mathcal{F})$ such that*

$$\mathbb{P}(A) = \int P(A) d\mu \tag{1}$$

for all $A \in \mathcal{F}$.[10] *The probability μ is called the mixing measure.*

---

[7] A standard result is that regular conditional probabilities exist whenever Ω is a Polish space and $\mathcal{F}$ is its Borel sigma-algebra. See, for instance, Durrett (2010).

[8] If $\mathcal{F}$ is countably generated—again, a pretty common feature—then regular conditional probabilities are unique up to an exceptional set of ω with $\mathbb{P}$ measure 0.

[9] One could also consolidate the terminology by viewing P as a probability-measure-valued measurable function. I have avoided this point of view in the main text in order to keep the exposition simple.

[10] Note that $P(A)$ denotes a $\mathcal{G}$-measurable function of Ω, and is shorthand for $P(A, \cdot)$.

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

The next condition is an extension of Mixing to bounded random variables. It says that initial expected values are mixtures of possible future expected values. It is a generalization to arbitrary probability spaces of what van Fraassen (1999) calls The General Reflection Principle.

DEFINITION 3. *Say that an updating policy P satisfies* The General Reflection Principle (GRP) *if and only if, for all bounded random variables f, there exists a probability measure μ_f on $(\Omega, \mathcal{F})$ such that*

$$\int f d\mathbb{P} = \int \int f dP d\mu_f. \tag{2}$$

The probability measure μ_f is called the mixing measure associated with f.

The left-hand side of (2) is the prior expected value of f. The right-hand side is a mixture of the possible future expected values of f, where the mixture is taken with respect to μ_f. Note that the order of quantification here is different from the order in the definition of Mixing. In the definition of GRP, the mixing measure is allowed to depend on f. Different random variables may have different mixing measures associated with them. In the definition of Mixing, however, the *same* mixing measure μ works for all events A; μ is not allowed to depend on A.

What can be said to motivate Mixing and GRP? Both conditions capture the idea that an agent's initial judgments (of probability, of expected value) ought to be constrained in certain ways by the judgments she plans to adopt after acquiring new information. To make sense of these conditions, it's helpful to consider what it means for them to fail. Take Mixing first. One way for Mixing to fail is if, say, $\mathbb{P}(A) = 0.5$ for some event A and $P(A, \omega) > 0.6$ for all $\omega \in \Omega$. To see that Mixing fails here, just recall what Mixing means when Ω is finite and note that no convex combination of numbers strictly greater than 0.6 can be equal to 0.5. To see what the failure of Mixing amounts to in this example, recall how we interpret the relevant objects. $\mathbb{P}(A)$ is the agent's initial probability judgment about A, and $P(A, \omega)$ is the probability for A that the agent plans to adopt upon gaining partial information about the true ω. In the example, then, the agent is in the following predicament. Initially, prior to gaining the partial information in $\mathcal{G}$, she plans to increase her probability for A *no matter what* $\mathcal{G}$ reveals about the true outcome. If the agent expects to be *better informed* after learning the partial information in $\mathcal{G}$, then this combination of attitudes seems odd. It seems that the agent ought to increase her current probability judgment above 0.6, since that is what she plans to do in her more informed future anyway. On the other hand, if the agent expects that the information represented by $\mathcal{G}$ is misleading or otherwise pathological, then this combination of attitudes might be reasonable. It is well known from the philosophical literature devoted to finding counterexamples to van Fraassen's (1984) reflection principle that not all changes of mind need to satisfy Mixing. If an agent expects to get drunk and become systematically overconfident about some of the events of interest to her, then she might reasonably violate Mixing. An idea that I am partial to is: if a change of mind is prompted by *genuine learning* of *valuable information*, then Mixing is a natural requirement. For then it is simply an injunction to defer to future states in which one has more valuable information.[11]

[11] The distinction between valuable information and mere information is an important one in epistemology (Levi, 1980). Although I won't go into details here, suffice it to say that it

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

So, I think that the *most* one can say in favor of Mixing is that it is characteristic of certain kinds of changes of mind—those that are the result of learning valuable information.[12] It is not, I think, a universal rationality requirement that applies to all pairs of current and future probability judgments.[13] By this I mean that there may be scenarios, like the one mentioned above, in which a rational agent has an initial probability $\mathbb{P}$ that is not a mixture of the probabilities P that she plans to adopt in the future. Perhaps rationality forbids this state of affairs when P is a policy for learning valuable information, but my aim here is not to argue these points at length. As we will see in the next section, the role that both Mixing and GRP play in this paper's argument is simply to delineate a class of learning scenarios in which Kolmogorov conditionalization is required.

Having discussed Mixing, what can be said about GRP? The points just raised about Mixing carry over *mutatis mutandis* to GRP, but, in the case of GRP, a bit more can be said. In particular, GRP has an important decision-theoretic interpretation. Let us view the random variables f that figure into the definition of GRP as *acts* in roughly the sense of Savage (1954). An act is a mapping from the uncertain outcome space Ω into a space of utility payoffs, which we assume can be measured with numbers in $\mathbb{R}$. The integral $\int f d\mathbb{P}$ then represents the agent's expected utility for f prior to learning, and $\int f dP$ is a policy for revising expected utilities in the light of the information provided by $\mathcal{G}$. Say that an act is acceptable with respect to some probability measure if it has non-negative expected utility according to that probability measure. Then, a consequence of GRP is, if f is acceptable to P *no matter what* $\mathcal{G}$ reveals about the true outcome, then f must be acceptable to $\mathbb{P}$ as well. Formally, if $\int f dP$ is uniformly greater than or equal to 0, then $\int f d\mathbb{P} \geq 0$. Under GRP, deference to the more informed future extends from probability judgments (as discussed above) to judgments about acceptable acts. If GRP is satisfied, then the collection of acts that the agent views as acceptable now is constrained by the collection of acts that she expects to view as acceptable after learning.[14]

As with Mixing, I think it is implausible to suggest that GRP is a universal rationality constraint that applies to all pairs of current and future probabilities. Continuing with the decision-theoretic interpretation, there may be scenarios in which it is reasonable for an agent to judge that an act is acceptable now and judge that the same act will not be acceptable come what may in the future. Perhaps the agent expects a medication she is about to take to radically alter her judgments about the acts she is considering. Again, one might be able to say that GRP is characteristic of scenarios in which changes of mind come about by learning valuable information. I find that idea plausible but won't defend it here. The next step in the present paper is to demonstrate a mathematical result, and then to discuss an argument for Kolmogorov conditionalization that is based on the result.

---

is arguably unreasonable to expect an agent to defer to more informed future states if the additional information possessed in those states is without value.

[12] For a discussion of similar issues, see Skyrms (1990a, 1990b, 1996) and Huttegger (2014).

[13] See Huttegger (2017) for a careful and detailed discussion of this view.

[14] In this formulation, it is clear that GRP is closely related to the Pareto conditions that play a role in some of the central results of social welfare theory (Harsanyi, 1955; Zhou, 1997).

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

§5. A characterization result and an argument. The characterization result that I establish in this section generalizes theorems T1 and T2 from van Fraassen (1999). After stating and proving the result, I will discuss how it gives rise to an argument for Kolmogorov conditionalization.

First, however, let me clarify exactly how the result to follow generalizes van Fraassen. In particular, I would like to clarify how the notion of a proper updating policy arises in van Fraasen's framework. The framework of this paper reduces to van Fraassen's framework by assuming that Ω is finite. In this case, every sigma-subalgebra $\mathcal{G}$ is generated by a finite partition of Ω. For simplicity, let us identify $\mathcal{G}$ with the partition that generates it and write $\mathcal{G} = \{G_1, \ldots, G_n\}$. Since P is $\mathcal{G}$-measurable in its second argument, $P(A, \cdot)$ is constant on the atoms of $\mathcal{G}$ for all events A. Let P_i be the probability measure attained by P on G_i, where i ranges between 1 and n. In the results that we generalize, van Fraassen assumes that $\{P_i : 1 \leq i \leq n\}$ is an "orthogonal family," by which he means $P_i(G_i) = 1$ for all i. This is equivalent to assuming that P is proper. Indeed, given the orthogonality assumption and any $\omega \in \Omega$ and $G_i \in \mathcal{G}$, let $\omega \in G_j$. Then, $P(G_i, \omega) = P_j(G_i)$ equals 1 if $i = j$ and 0 otherwise; or, in other words, equals 1 if $\omega \in G_i$ and 0 otherwise; which is to say, equals $1_{G_i}(\omega)$. Conversely, assuming that P is proper, we have $P(G_i, \omega) = 1_{G_i}(\omega)$ for all ω and all i; or, in other words, for all i, j and all $\omega \in G_j$, we have that $P_j(G_i)$ equals 1 if $i = j$ and 0 otherwise; which is to say, $P_i(G_i) = 1$ for all i. Thus, we see that properness is the appropriate generalization of van Fraassen's orthogonality concept. The rest of the conditions that van Fraassen employs in his results are easily seen to be special cases of Mixing and GRP. As I pointed out above, when Ω is finite, both Mixing and GRP can be formulated using the language of spanning intervals and convex combinations.

With that clarification out of the way, I can now state my result. Simply put, it says that in proper Kolmogorov learning scenarios—those in which the agent actually gains full membership knowledge about $\mathcal{G}$—the conditions that we introduced above—Kolmogorov conditionalization, Mixing, and GRP—are equivalent.

THEOREM 1. *Let $(\Omega, \mathcal{F}, \mathcal{G}, \mathbb{P}, P)$ be a proper Kolmogorov learning scenario in which a regular conditional probability $\mathbb{P}_{\mathcal{G}}$ for $\mathbb{P}$ given $\mathcal{G}$ exists. Then, the following assertions are equivalent*:

(I) *P is given by Kolmogorov conditionalization.*[15]
(II) *P satisfies Mixing.*
(III) *P satisfies The General Reflection Principle.*

Proof. That (I) implies (II) is immediate from (iii); in the definition of Mixing, we can let $\mu = \mathbb{P}$.

That (II) implies (III) follows from a straightforward monotone class argument. We will omit most of the details. Assuming that (II) holds, let μ be the mixing measure given by (1). Let $\mathcal{C}$ be the collection of random variables f for which (2) holds with $\mu_f = \mu$. Then, by (II), $\mathcal{C}$ contains the indicator function for every event in $\mathcal{F}$. It is now easy to show, using the linearity of the integral and the monotone convergence theorem, that $\mathcal{C}$ is closed under addition, scalar multiplication, and bounded monotone limits. It follows that $\mathcal{C}$ contains every bounded random variable, so (III) holds. (In

[15] Note that by Theorem 1 in Blackwell & Dubins (1975), the existence of a proper regular conditional probability implies that $\mathcal{G}$ is countably generated.

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

fact, we showed something *prima facie* stronger than (III), since we established that μ satisfies (2) for all f. The interesting part of the proof is the demonstration that, in the presence of (iv), the weak-looking (III) is strong enough to imply (I).)

We conclude by establishing that (III) implies (I). Assume that (III) holds. In order to show that P is a version of $\mathbb{P}_{\mathcal{G}}$, it suffices to establish (iii). To that end, let $A \in \mathcal{F}$ and $G \in \mathcal{G}$ be given. Write $f = 1_G P(A \cap G) - 1_A 1_G$, so that f is a bounded random variable. Now, P satisfies (iii) if and only if $\int f d\mathbb{P} = 0$. To see this, observe that

$$\begin{aligned}\int f d\mathbb{P} = 0 &\Leftrightarrow \int_G P(A \cap G) d\mathbb{P} = \mathbb{P}(A \cap G) \\ &\Leftrightarrow \int_G P(A) d\mathbb{P} = \mathbb{P}(A \cap G),\end{aligned}$$

where, in passing to the second line, we have used the fact that P is proper: for all $\omega \in G$, $P(G, \omega) = 1$ implies $1_G P(A \cap G) = 1_G P(A)$. To show that $\int f d\mathbb{P} = 0$, let $P_\omega = P(\cdot, \omega)$ and observe that, for all $\omega \in \Omega$,

$$\int f dP_\omega = \int_G P(A \cap G) dP_\omega - P_\omega(A \cap G).$$

The function being integrated on the right-hand side of the last equation, namely $1_G P(A \cap G)$, is $\mathcal{G}$-measurable. Thus, the integral $\int_G P(A \cap G) dP_\omega$ can be evaluated with respect to the restriction of P_ω to $(\Omega, \mathcal{G})$. By the fact that P is proper, the restriction of P_ω to $(\Omega, \mathcal{G})$ is point mass at ω, which we denote by δ_ω. Hence,

$$\int f dP_\omega = \int_G P(A \cap G) d\delta_\omega - P_\omega(A \cap G) = 1_G(\omega) P_\omega(A \cap G) - P_\omega(A \cap G).$$

There are now two cases to consider in order to evaluate the right-most formula. If $\omega \in G$, then the formula reduces to $P_\omega(A \cap G) - P_\omega(A \cap G) = 0$; if $\omega \notin G$, then, because P is proper, $P_\omega(A \cap G) = 0$ and the formula vanishes again. We have shown that $\int f dP_\omega = 0$ for all $\omega \in \Omega$; the function $\int f dP$ is everywhere 0. Finally, let μ_f be the mixing measure associated with f in (2), and use (III) to get

$$\int f d\mathbb{P} = \int \int f dP d\mu_f = \int 0 d\mu_f = 0.$$

As shown above, this suffices to establish (I) and concludes the proof. □

So much for the result. What does it tell us about Kolmogorov conditionalization as a policy for updating probability judgments? The result gives rise to a simple argument for Kolmogorov conditionalization by showing that, in a particular class of learning scenarios, Kolmogorov conditionalization is the only correct updating rule for an agent to adopt. What learning scenarios are those? They are the ones in which the agent expects to gain membership knowledge about a sigma-subalgebra of the events of interest to her and in which her initial probability judgments are constrained by her updating policy in the manner prescribed by The General Reflection Principle (or Mixing). In such scenarios, Theorem 1 tells us that the only updating policies that an agent can consistently adopt are those given by Kolmogorov conditionalization. On pain of inconsistency, then, an agent should adopt an updating policy given by Kolmogorov conditionalization in any proper learning scenario in which she expects GRP to hold.

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

As van Fraassen has stressed, this argument for conditionalization is not universal, since it shows that conditionalization is correct only under "certain special conditions" (1999, p. 93). As I remarked above, it is implausible to expect all learning scenarios to be proper. Scenarios in which agents gain "uncertain evidence"—the kinds of scenarios with which Richard Jeffrey's radical probabilism is concerned and for which he developed probability kinematics—are improper. Moreover, I have suggested that it is implausible to view GRP (and Mixing) as a universal rationality constraint that applies to all pairs of current and future probabilities. It is important to note, then, that the present argument does *not* show that all reasonable changes of mind are given by conditionalization. For there may well be scenarios in which an agent reasonably violates GRP. Perhaps the argument for Kolmogorov conditionalization can be extended along the lines suggested in §4. If GRP is characteristic of genuine learning scenarios in which valuable information is gained, then our argument shows that Kolmogorov conditionalization is the only correct updating policy in learning scenarios where the agent expects to gain valuable membership knowledge about $\mathcal{G}$. The antecedent of that conditional is in need of defense, however—a task I have left for another time. Van Fraassen has suggested a different extension of the argument by suggesting that properness and GRP are characteristic of well-designed scientific experiments. If that's right, then the argument shows that Kolmogorov conditionalization is the only correct updating policy for a scientist to adopt, unless she suspects that her experiments are poorly designed.

Contrary to the claims that one sometimes hears about conditionalization being the only rational policy for updating probability judgments, the present argument for conditionalization is circumscribed by properness and GRP. When either of those conditions fails to obtain, the argument is silent. The reader should be reminded, however, that there are *other* arguments supporting conditionalization that do not rely on properness and GRP. As we have seen, Rescorla's Dutch book theorems apply in improper learning scenarios. Although Rescorla himself is reluctant to use his theorems to argue in favor of Kolmogorov conditionalization,[16] it is easy to see how such an argument would go. Insofar as an agent ought to be concerned with avoiding Dutch books in Kolmogorov learning scenarios, the only permissible updating policies for her to adopt are Kolmogorov conditionalizations. For those who find Dutch book arguments dubious (for a survey, see Hájek, 2009), the argument in the present paper offers different considerations in favor of Kolmogorov conditionalization.

Before concluding the paper, I will provide a few more mathematical results and remarks that extend what was shown in Theorem 1.

§6. Additional mathematical remarks. The main result of the paper, Theorem 1, is stated in a way that makes it a straightforward generalization of the results of van Fraassen (1999). It's not too difficult, however, to push the argument a bit farther. In this section, I present some additional conditions that are equivalent to Kolmogorov conditionalization in proper Kolmogorov learning scenarios.

The first condition, like GRP and Mixing, formalizes a kind of deference to future probability judgments. It says that if the potential future probabilities associated with an updating policy unanimously assign a bounded random variable an expected value

[16] See Rescorla (2018, pp. 708–709).

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

of zero, then the expected value of the random variable must be zero with respect to the initial probability as well. Before stating the definition, recall the notation introduced in the proof of Theorem 1: $P_\omega = P(\cdot, \omega)$.

DEFINITION 4. *Say that an updating policy P is Pareto if and only if the following condition holds for all bounded random variables f*:

$$\text{If } \int f\,dP_\omega = 0 \text{ for all } \omega \in \Omega, \text{ then } \int f\,d\mathbb{P} = 0. \tag{3}$$

The name of the condition is meant to connect it with the discussion of GRP in §4 (see, in particular, footnote 14).

The final two conditions weaken some of the countable additivity assumptions that we made above by allowing mixing measures to be merely finitely additive (we continue to assume that $\mathbb{P}$ and P_ω are countably additive).

DEFINITION 5. *Say that an updating policy P satisfies* FA Mixing *if and only if there exists a finitely additive probability measure μ on $(\Omega, \mathcal{F})$ such that*

$$\mathbb{P}(A) = \int P(A)d\mu \tag{4}$$

for all $A \in \mathcal{F}$. The probability μ is called the finitely additive mixing measure.

DEFINITION 6. *Say that an updating policy P satisfies* The Finitely Additive General Reflection Principle (FAGRP) *if and only if, for all bounded random variables f, there exists a finitely additive probability measure μ_f on $(\Omega, \mathcal{F})$ such that*[17]

$$\int f\,d\mathbb{P} = \int\int f\,dP d\mu_f. \tag{5}$$

The probability measure μ_f is called the finitely additive mixing measure associated with f.

In finite probability spaces, where every finitely additive probability measure is countably additive as well, Mixing is equivalent to FA Mixing, and GRP is equivalent to FAGRP. In general, however, the latter conditions are weaker.

We now show that the Pareto condition, FA Mixing, and FAGRP are equivalent to Kolmogorov conditionalization, Mixing, and GRP in proper Kolmogorov learning scenarios.

THEOREM 2. *Let $(\Omega, \mathcal{F}, \mathcal{G}, \mathbb{P}, P)$ be a proper Kolmogorov learning scenario in which a regular conditional probability $\mathbb{P}_\mathcal{G}$ for $\mathbb{P}$ given $\mathcal{G}$ exists. Then, the following assertions are equivalent to the assertions (I), (II), and (III) in the statement of Theorem 1*:

(IV) *P satisfies FA Mixing.*
(V) *P satisfies FAGRP.*
(VI) *P is Pareto.*

Proof. The implication from (II) to (IV) is immediate.

(V) follows from (IV) by an argument like the one showing that (II) implies (III) in the proof of Theorem 1, using the fact that the finitely additive integral is linear

[17] For a brief overview of finitely additive integration, see Aliprantis & Border (2006). For a detailed treatment, see Rao & Rao (1983).

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

and the fact that every bounded random variable can be uniformly approximated by a sequence of simple random variables.

To show that (V) implies (VI), assume that P satisfies FAGRP and that $\int f dP$ is everywhere 0. Then, with μ_f the finitely additive mixing measure associated with f, we have, by (5),

$$\int f d\mathbb{P} = \int \int f dP d\mu_f = \int 0 d\mu_f = 0.$$

So, P is Pareto.

Finally, to show that (VI) implies (I), assume that P is Pareto. Repeat the proof that (III) implies (I) in Theorem 1 until the point that $\int f dP$ is observed to be everywhere 0 and use the Pareto condition to conclude immediately. □

Although none of the steps in the proofs above was especially difficult, it is interesting to conclude with the result that, if P is proper, then it is a version of the regular conditional probability $\mathbb{P}_\mathcal{G}$ just in case it satisfies FA Mixing. Blackwell & Dubins (1975, pp. 741–742) have observed that for proper P, (iii) is implied by the weaker condition

$$\mathbb{P}(A) = \int P(A) d\mathbb{P} \text{ for all } A \in \mathcal{F}. \tag{6}$$

So, if P is proper, then to verify that P is a version of $\mathbb{P}_\mathcal{G}$ one can check just (6) instead of the stronger (iii). One upshot of the foregoing results is that the weak condition (6) that appears in the Blackwell-Dubins observation can be weakened even more, by replacing $\mathbb{P}$ on the right-hand side with any finitely additive probability measure.

§7. Conclusion. Kolmogorov conditionalization is the only correct policy for updating one's probability judgments in proper learning scenarios in which one expects The General Reflection Principle to hold. The argument showing this and the theorem on which it is based are generalizations of arguments and results due to van Fraassen. What we have established, essentially, is that van Fraassen's main points do not depend on assumptions that make learning scenarios *elementary*: we needn't assume that outcome spaces are finite, that information is represented by a finite partition, and so on.

I take the broader significance of the paper—and the research program of which it is a part (see §1)—to be twofold. On the one hand, it is important that Bayesian epistemology, and epistemology generally, not be divorced from scientific practice. Probabilisitic learning is not always elementary, and our epistemology should be able to account for that. The present paper, together with papers mentioned above, shows that this can be done. Philosophical arguments supporting elementary conditionalization can be extended to account for the most widely used mathematical generalization of it: Kolmogorov conditionalization. On the other hand, this work also shows that the relevant philosophical concepts—invulnerability to Dutch book, The General Reflection Principle, and so on—are durable and will find applications in scenarios that require nonelementary probability models.

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

BIBLIOGRAPHY

Aliprantis, C. D., & Border, K. C. (2006). *Infinite Dimensional Analysis*: *A Hitchhiker's Guide*. Berlin: Springer-Verlag.

Blackwell, D., & Dubins, L. E. (1975). On existence and non-existence of proper, regular, conditional distributions. *The Annals of Probability*, **3**(5), 741–752.

Blackwell, D., & Ryll-Nardzewski, C. (1963). Non-existence of everywhere proper conditional distributions. *Annals of Mathematical Statistics*, **34**(1), 223–235.

de Finetti, B. (1974). *Theory of Probability*, Vol. 1. Chichester: John Wiley & Sons.

Dubins, L. E. (1975). Finitely additive conditional probabilities, conglomerability and disintegrations. *The Annals of Probability*, **3**(1), 89–99.

Durrett, R. (2010). *Probability*: *Theory and Examples*. New York: Cambridge University Press.

Good, I. J. (1967). On the principle of total evidence. *The British Journal for the Philosophy of Science*, **17**(4), 319–321.

Greaves, H., & Wallace, D. (2006). Justifying conditionalization: conditionalization maximizes expected epistemic utility. *Mind*, **115**(459), 607–632.

Gyenis, Z., Hofer-Szabo, G., & Rédei, M. (2017). Conditioning using conditional expectations: the Borel–Kolmogorov paradox. *Synthese*, **194**(7), 2595–2630.

Gyenis, Z., & Rédei, M. (2017). General properties of Bayesian learning as statistical inference determined by conditional expectations. *The Review of Symbolic Logic*, **10**(4), 719–755.

Hájek, A. (2009). Dutch book arguments. In Anand, P., Pattanaik, P., & Puppe, C., editors. *The Handbook of Rationality and Social Choice*. Oxford: Oxford University Press, pp. 173–195.

Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, **63**(4), 309–321.

Huttegger, S. M. (2013). In defense of reflection. *Philosophy of Science*, **80**(3), 413–433.

———(2014). Learning experiences and the value of knowledge. *Philosophical Studies*, **171**(2), 279–288.

———(2015). Merging of opinions and probability kinematics. *The Review of Symbolic Logic*, **8**(4), 611–648.

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

Huttegger, S. M., & Nielsen, M. (2019). Generalized learning and conditional expectation. *Philosophy of Science*, forthcoming.

Jeffrey, R. (1965). *The Logic of Decision*. Chicago: University of Chicago Press.

———(1992). *Probability and the Art of Judgment*. Cambridge: Cambridge University Press.

Kolmogorov, A. N. (1950). *Foundations of the Theory of Probability*. New York: Chelsea Publishing Company.

Leitgeb, H., & Pettigrew, R. (2010a). An objective justification of bayesianism i: measuring inaccuracy. *Philosophy of Science*, **77**(2), 201–235.

———(2010b). An objective justification of bayesianism ii: the consequences of minimizing inaccuracy. *Philosophy of Science*, **77**(2), 236–272.

Levi, I. (1980). *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.

Ramsey, F. P. (1931). Truth and probability. In Braithwaite, R. B., editor. *The Foundations of Mathematics and Other Essays*. New York: Kegan, Paul, Trench, Trubner, & Co., pp. 156–198.

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press

Rao, K. B., & Rao, M. B. (1983). *Theory of Charges*: *A Study of Finitely Additive Measures*. Pure and Applied Mathematics, Vol. 109. London: Academic Press.
Rescorla, M. (2018). A Dutch book theorem and converse Dutch book theorem for Kolmogorov conditionalization. *The Review of Symbolic Logic*, **11**(4), 705–735.
Savage, L. J. (1954). *The Foundations of Statistics*. New York: Wiley.
Seidenfeld, T. (2001). Remarks on the theory of conditional probability: Some issues of finite versus countable additivity. In Hendricks, V. F., Pedersen, S. A., & Jørgensen, K. F., editors. *Probability Theory*. Dordrecht, The Netherlands: Klewer, pp. 167–178.
Seidenfeld, T., Schervish, M. J., & Kadane, J. B. (2001). Improper regular conditional distributions. *The Annals of Probability*, **29**(4), 1612–1624.
Skyrms, B. (1990a). *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.
———(1990b). The value of knowledge. *Minnesota Studies in the Philosophy of Science*, **14**, 245–266.
———(1996). The structure of radical probabilism. *Erkenntnis*, **45**(2–3), 285–297.
Teller, P. (1973). Conditionalization and observation. *Synthese*, **26**(2), 218–258.
van Fraassen, B. C. (1984). Belief and the will. *The Journal of Philosophy*, **81**(5), 235–256.
———(1989). *Laws and Symmetry*. Oxford: Clarendon Press.
———(1999). Conditionalization, a new argument for. *Topoi*, **18**(2), 93–96.
Zhou, L. (1997). Harsanyi's utilitarianism theorems: general societies. *Journal of Economic Theory*, **72**(1), 198–207.

DEPARTMENT OF PHILOSOPHY
THE UNIVERSITY OF SYDNEY
NSW 2006, AUSTRALIA
E-mail: michael.nielsen@sydney.edu.au

https://doi.org/10.1017/S1755020320000179 Published online by Cambridge University Press