In Defense of Reflection
Author(s): Simon M. Huttegger
Source: *Philosophy of Science*, Vol. 80, No. 3 (July 2013), pp. 413-433
Published by: The University of Chicago Press on behalf of the Philosophy of Science Association
Stable URL: http://www.jstor.org/stable/10.1086/671427
Accessed: 29/06/2013 11:00

In Defense of Reflection

Simon M. Huttegger*†

I discuss two ways of justifying reflection principles. First, I propose that an undogmatic reading of dynamic Dutch book arguments provides them with a sound foundation. Second, I show also that minimizing expected inaccuracy leads to a novel argument for reflection principles. The required inaccuracy measures comprise a natural class of functions that can be derived from a generalization of a condition known as propriety or immodesty. This shows that reflection principles are an essential feature not just of consistent degrees of belief but also of degrees of belief that approximate truth.

> And then, as he pushed through a hedge into a field untended, there suddenly close before him in the field was, as his father had told, the frontier of twilight. It stretched across the fields in front of him, blue and dense like water; and things seen through it seemed misshapen and shining. (Lord Dunsany, *The King of Elfland's Daughter*)

1. Introduction. Reflection principles relate one's anticipated future opinions to one's current opinions. One way to couch reflection is to say that my "current opinion about event E must lie in the range spanned by the possible opinions I may come to have about E at later time t, as far as my present opinion is concerned" (van Fraassen 1995, 16). In the theory of arbitrage, the fundamental theorem of asset pricing gives an exact statement of this idea (Skyrms 2006). If the opinions involved are precise probabilities, the fore-

Received March 2013; revised May 2013.

*To contact the author, please write to: Department of Logic and Philosophy of Science, University of California, Irvine, Social Science Plaza A, Irvine, CA 92697; e-mail: shuttegg@uci.edu.

†For helpful comments on earlier drafts of this article, I want to thank Kenny Easwaran, Brian Skyrms, Bas van Fraassen, Jonathan Weisberg, Kevin Zollman, and an anonymous referee for this journal.

Philosophy of Science, 80 (July 2013) pp. 413–433. 0031-8248/2013/8003-0009$10.00

going reflection principle reads as follows:

(**R1**) An agent's current belief $\mathbb{P}[A]$ in event A should lie in the interval spanned by any set of values $\mathbb{P}_f[A]$ of her anticipated future degree of belief in A that has probability 1.[1]

Two further principles go by the name reflection:

(**R2**) An agent's current belief $\mathbb{P}[A]$ in event A should be the expectation of her anticipated future degrees of belief $\mathbb{P}_f[A]$.
(**R3**) An agent's current degree of belief in event A given that her anticipated future degree of belief $\mathbb{P}_f[A] = r$ should be equal to r with probability 1, whenever the event $\mathbb{P}_f[A] = r$ has positive probability.

Similar statements can be formulated for quantities other than future degrees of belief. This, as well as what I mean by 'anticipated future degrees of beliefs', is explained in section 2. The main purpose of section 2 is to introduce the general concepts of *conditional expectation* and *conditional probability* and to note in what sense they imply R1 and R2. This can be used to argue that if future degrees of belief are given by conditional probabilities, then R1 and R2 hold. The burden of proof therefore lies on showing that one's anticipated future degrees of belief should be equal to conditional probabilities. This leads, among other things, to the question of justifying principles like R3.

Dutch book arguments are one way to show that one's future degrees of belief should be given by a conditional probability. They can also be used to defend a principle such as R3 (see Goldstein 1983; van Fraassen 1984). This approach has drawn rather fierce criticisms on itself and on reflection principles (for discussions and criticisms, see Levi 1987; Christensen 1991; Talbott 1991; Maher 1992; Bacchus, Kyburg, and Thalos 1995; Bovens 1995; Arntzenius 2003; Briggs 2009). I argue in section 3 that a liberal reading of Dutch book arguments helps in understanding the role of reflection for belief change and puts into context various aspects of the counterexamples to reflection. Reflection, properly understood, turns out to be a requirement of epistemic rationality for any agent who considers herself to update beliefs rationally.

The main part of my article develops a new justification for reflection principles in terms of expected accuracy. In sections 4 and 5, I show that con-

1. The probability 1 qualification is important here and in R3, whenever a probability space is infinite. It can be disregarded in finite probability spaces, where all events have positive probability. This point is sometimes ignored in the philosophical literature.

ditional expectations minimize expected inaccuracy for a natural class of inaccuracy measures that are based on the well-known condition of propriety or immodesty. As a consequence of this result, one can derive generalizations of the reflection principles R1–R3. Approaches based on minimizing expected inaccuracy have recently been used by Greaves and Wallace (2006), Leitgeb and Pettigrew (2010a), and Easwaran (2013) to justify various aspects of Bayesianism. These approaches are in the philosophical tradition of so-called nonpragmatic vindications of probabilism (Joyce 1998), with the distinguishing feature that accuracy is weighed by one's own prior probabilities.

I would like to emphasize that it is not my goal to show that reflection principles are requirements for rational belief change under all circumstances. Rather, the aim of this article is to explicate both the framework and the assumptions that are required for deriving reflection principles. This will lead, I hope, to some clarity as to what those of us mean who claim that reflection is a requirement for rational belief change.

2. Conditional Probability and Reflection. I start with a rather careful introduction of the general concepts of conditional probability and conditional expectation. These concepts play a crucial role in understanding reflection principles.

The easiest case is one in which you perform an experiment with finitely many possible results $\{E_i\}$ that are assumed to form a partition. If you update your probabilities by conditioning on the outcome of the experiment, the expected value of your new probabilities is given by

$$\mathbb{P}[A] = \sum_i \mathbb{P}[A|E_i]\mathbb{P}[E_i], \tag{1}$$

where the sum extends over all i such that $\mathbb{P}[E_i] > 0$. Your current probability $\mathbb{P}[A]$ is therefore given by the expectation of your conditional probabilities.

This is a very elementary observation. Two questions immediately suggest themselves. First, does this result hold for general probability spaces? And, second, is it necessary that the experimental outcomes form a partition?

To answer these questions, suppose that $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space, where Ω is a set of elementary events or—if you wish—possible worlds, and $\mathcal{F}$ is a σ-algebra of subsets of Ω. A σ-algebra is just an algebra that is also closed under countable unions and intersections. The members of $\mathcal{F}$ are called events or propositions. We assume that $\mathbb{P}$ is a countably additive probability measure on $(\Omega, \mathcal{F})$.

In what follows, $(\Omega, \mathcal{F}, \mathbb{P})$ is the subjective probability space of an agent. This means that Ω represents the worlds that the agent considers possible, the events in $\mathcal{F}$ are the propositions the agent can express and for which she has probabilities, and $\mathbb{P}$ is the agent's subjective probability measure. This is tantamount to saying that the agent is synchronically σ-coherent (see Adams 1962).

A $\mathcal{F}$-measurable random variable X is a function from Ω to the real numbers $\mathbb{R}$ such that every event $X^{-1}(B)$ is in $\mathcal{F}$ for each open subset B of $\mathbb{R}$. Loosely speaking, the events described by $\mathcal{F}$-measurable random variables do not go beyond what is expressible in $\mathcal{F}$.

If A is in $\mathcal{F}$, we denote the expectation of X on A by

$$E[X;\ A] = \int_A X d\mathbb{P}.$$

We can now define *conditional expectations*.[2] Suppose that X is a $\mathcal{F}$-measurable random variable and that $\mathcal{G}$ is a sub-σ-algebra of $\mathcal{F}$.[3] The conditional expectation $E[X|\mathcal{G}]$ of X given $\mathcal{G}$ is a $\mathcal{G}$-measurable random variable for which it is true that

$$E[E[X|\mathcal{G}]; G] = E[X; G] \text{ for all } G \text{ in } \mathcal{G}; \tag{2}$$

that is, on each G in $\mathcal{G}$, the random variables X and $E[X|\mathcal{G}]$ have the same expectation. It can be shown that $E[X|\mathcal{G}]$ exists and is almost surely unique (any two random variables for which [2] holds are equal up to a set of probability 0; see, e.g., Williams [1991] for details).

The general concept of *conditional probability* is obtained as a special case, by setting $X = I_A$ (the indicator of A).[4] In this case, $E[X|\mathcal{G}] = E[I_A|\mathcal{G}] = \mathbb{P}[A|\mathcal{G}]$. Hence, the conditional probability $\mathbb{P}[A|\mathcal{G}]$ is also a $\mathcal{G}$-measurable random variable.

The random variables $P[A|\mathcal{G}]$ and $E[X|\mathcal{G}]$ often reduce to their well-known counterparts in finite probability spaces. Suppose that $G \in \mathcal{G}$ is an atom of $\mathcal{G}$.[5] Then $\mathbb{P}[A|\mathcal{G}](\omega) = \mathbb{P}[A|G]$ for almost every ω in G. This shows that the general concept of conditional probability almost surely agrees with the standard ratio definition of conditional probability in this case.

2. The expectation of X on A should not be confused with the expectation of X given the set A, which is a special case of conditional expectation. The expectation of X on A is generally equal to the product of that conditional expectation with the probability of A.

3. That is, $\mathcal{G}$ is a σ-algebra, and $\mathcal{G} \subset \mathcal{F}$.

4. The indicator of A can be thought of as the truth value of A in world ω: $I_A(\omega) = 1$ if $\omega \in A$, and $I_A(\omega) = 0$ otherwise.

5. That is, $\mathbb{P}[G] > 0$ and $\mathbb{P}[E] = 0$ or $\mathbb{P}[G - E] = 0$, for every $E \in \mathcal{G}$ that is a subset of G.

There are generalizations of (1) for conditional expectations and conditional probabilities. If we set $G = \Omega$ in (2), we get

$$E[E[X|\mathcal{G}]] = E[X], \tag{3}$$

and this implies

$$E[\mathbb{P}[A|\mathcal{G}]] = \mathbb{P}[A]. \tag{4}$$

Equations (3) and (4) hold in general probability spaces. We can get back to a simple case like (1) by taking $\mathcal{G}$ to be generated by a finite partition $\{E_i\}$ of Ω; then it is true that

$$E[\mathbb{P}[A|\mathcal{G}]] = \sum_i \mathbb{P}[A|E_i]\mathbb{P}[E_i],$$

where the sum ranges over all i such that E_i has positive probability. (This follows since any such E_i is an atom of $\mathcal{G}$.) Hence, (4) is a proper generalization of the law of total probability, and (3) is its extension to random variables.

Reflection enters the picture by observing that (4) looks very similar to R2 and that something like R1 follows from this because of standard properties of expectation. What needs to be clarified, however, is in what sense the conditional probability $\mathbb{P}[A|\mathcal{G}]$ can capture the agent's anticipated future degrees of belief referred to in R1 and R2.[6]

There seem to be two possible interpretations. One is to view $\mathbb{P}[A|\mathcal{G}]$ as a *plan* to update one's degree of belief for A after being informed which member of $\mathcal{G}$ is true. This approach is used by Easwaran (2013). On this reading, $\mathcal{G}$ is viewed as describing the outcomes of an experiment that provides the agent with a new piece of information. A similar approach is taken by Greaves and Wallace (2006), who consider *acts* instead of plans. An act is a probability distribution chosen by the agent in response to receiving some piece of information.[7] More generally, van Fraassen (1995) talks about *policies* for opinion change. This also includes conditional expectations $E[X|\mathcal{G}]$ and the corresponding reflection principles.

6. This is pointed out by Weisberg (2007).

7. I agree with Easwaran that for general probability spaces, the concept of a plan appears to be superior: $\mathbb{P}[A|\mathcal{G}]$ is a random variable whose existence is guaranteed. But it need not always be possible to have a function $\mathbb{P}[\cdot|\mathcal{G}]$ that is a probability measure almost surely. This problem leads to the question of the existence of regular conditional probabilities, on which I say a bit more in sec. 4. For now, let me just say that the probability distributions Greaves and Wallace use as acts need not exist in infinite probability spaces. Both acts and plans can be viewed as dispositions to update in a particular way. See Greaves and Wallace (2006) and Easwaran (2013) for more on this issue.

Understanding anticipated future degrees of belief in terms of plans gives rise to a notion of reflection that Easwaran aptly calls *plan reflection*. The corresponding principles result if we substitute for 'anticipated future degree of belief' the 'degree of belief an agent plans to have' in R1 and R2.[8]

On the second interpretation, the agent is assumed to believe with probability 1 that she will update by conditioning on $\mathcal{G}$. In this case, $\mathbb{P}[A|\mathcal{G}]$ will be her future degree of belief for A in almost every world ω. This interpretation seems to be used by Weisberg (2007) and Briggs (2009). The assumption that the agent believes with probability 1 that she will condition may of course fail. But if it holds, her anticipated degrees of belief are given by her conditional probabilities, and principles R1 and R2 again follow.

I am not going to adjudicate between these two interpretations because I think that both are reasonable ways to make precise the notion of anticipated future degrees of belief. Let me point out, however, that neither of them allows for a truly diachronic notion of belief change. Both conceptualize the future beliefs of an agent as something she contemplates from her current point of view. This view of belief change is fundamentally synchronic. It is not inconsistent with, but seems to be at odds with, the way many authors understand conditionalization, namely, as a norm for belief change the agent ought to apply as she actually learns a proposition for certain and moves to a new probability measure.

I will not offer a full-fledged defense of why a synchronic reading of belief change is appropriate. I think there are some reasonable arguments in favor of such a view, and I will mention one in the context of Dutch books; for some others, see Easwaran's discussion of the difficulties of transferring synchronic norms for conditional probabilities to a truly diachronic setting (Easwaran 2013). What is more important at this point is to notice that taking a synchronic view of belief change seems to be an intrinsic part of the measure-theoretic view of conditional probability as a random variable that lives within an agent's current probability space.

Thus, we see that the reflection principles R1 and R2 hold for conditional probabilities, provided that they are interpreted in a certain way. Before discussing principle R3, I would like to put emphasis on two features that are needed for (3) and (4) to count as reflection principles. In the first place, X is assumed to be $\mathcal{F}$-measurable; this excludes, for example, the possibility that an agent considers conditional probabilities for events $B \subset \Omega$ that are not members of $\mathcal{F}$. Moreover, $\mathcal{G}$ needs to be a subset of $\mathcal{F}$. The agent does not consider updating on information that she cannot currently express.[9] If

8. The resulting principles are, I think, those that van Fraassen (1995) has in mind when he says that conditionalization implies reflection.

9. This excludes Sleeping-Beauty-like cases.

these two assumptions do not hold, conditional expectations as introduced above are not defined.

Let me now turn to principle R3. One important point about conditional expectation is that the σ-algebra $\mathcal{G}$ can describe a variety of situations. An example that was already mentioned is the situation in which $\mathcal{G}$ is generated by the outcomes of an experiment. But $\mathcal{G}$ might also be generated by the agent's anticipated future degrees of belief, in the following way. Suppose that the agent finds herself in a highly unstructured learning situation. There is no nontrivial partition or σ-algebra that could serve as a basis for updating her beliefs. Perhaps the agent expects an unexpected informational input, or she cannot describe the events that she learns precisely enough in her language, or maybe she is just going to think about some topic. This kind of situation is called black-box learning in Skyrms (1990).

More formally, consider a random variable Y for an agent's future degree of belief regarding some fixed event A in $\mathcal{F}$. The value $Y(\omega)$ is her anticipated future degree of belief for A, if ω is the true state of the world, and Y captures the black-box nature of the learning situation since nothing whatsoever is assumed about the structure of the agent's learning experience. As in the case of conditional probability, Y may be understood in two ways: as a plan to update one's beliefs in response to the learning experience or as the quantity of which an agent believes, with probability 1, that it will be her degree of belief.

In order for Y to be well-defined, we need to assume that each ω contains information about the agent's degree of belief in A at world ω. The set Ω must therefore be quite rich. If, for example, the agent considers all real numbers in the interval $[0, 1]$ to be possible future degrees of belief for A, then the worlds in Ω need to reflect all these possible cases. We can also generalize from degrees of belief to estimates of unknown quantities. If X is a random variable, we can let Y be the (currently unknown) future best estimate of X.

Provided that Y is well-defined, let $\sigma(Y)$ be the smallest σ-algebra making Y measurable.[10] The σ-algebra $\sigma(Y)$ contains all propositions that can be described by Y. We assume that $\sigma(Y)$ is a subset of $\mathcal{F}$. This means that the agent already grasps all events that can be expressed with Y.

We thus require both that Y is well-defined and that $\sigma(Y)$ is a subset of $\mathcal{F}$. I do not take issue with these two assumptions for Y here. I only want to make them explicit and also draw attention to the fact that they are needed for conditional expectations. If $E[X|Y]$ denotes the conditional expectation of X given $\sigma(Y)$, then $E[X|Y]$ and $\mathbb{P}[A|Y]$ exist provided that our two assumptions are met.

10. The set $\sigma(Y)$ is the smallest σ-algebra that contains all sets of the form $Y^{-1}(B)$, where B is an open subset of $\mathbb{R}$.

Let's return to R3. In the present context, this principle says that, with probability 1, $Y = \mathbb{P}[A|Y]$; in fact, this is the precise formulation of R3 for general probability spaces.[11] The corresponding formula for random variables X is $Y = E[X|Y]$. This relation is characteristic for *martingales*. The reflection principle R3 is, in general, nothing but the martingale property. If we view Y as the currently unknown future best estimate of the true value X of some quantity, then it makes sense that $E[X|Y] = Y$, if we note that $E[X]$ can be viewed as the current best estimate of X.

The martingale property is not something that follows from any of the results that were mentioned so far. The approach I am going to focus on derives the martingale property within the framework of minimizing expected inaccuracy. But first, I briefly discuss the standard approach, which proceeds in terms of dynamic coherence.

3. The Role of Dutch Books. Since martingales are sequences of fair gambles, it is quite clear that dynamic coherence can be used to justify the martingale property. One dynamic coherence argument for reflection is due to van Fraassen (1984). It is formally the same as David Lewis's Dutch book argument for conditionalization (see Teller 1973). Goldstein (1983) puts forward a more general coherence argument that also applies to conditional expectations.

I will not repeat any of these arguments in detail here since they are well known. Let me just say that the general structure of a dynamic coherence argument is very similar to a synchronic coherence argument. An agent announces her fair-betting odds concerning events in $\mathcal{F}$, the basic idea being that the agent can choose the betting odds for a proposition but cannot choose the side of the bet.[12] Fair-betting odds are supposed to measure an agent's beliefs. By a suitable normalization, they are mapped to the real numbers. The resulting numbers are called the agent's degrees of beliefs for those events involved in the bets. A Dutch book argument then shows that if these numbers do not obey the laws of the probability calculus, the agent will end up with a net loss from a set of bets that she individually deems fair.

This line of reasoning can be applied to conditionalization (Teller 1973), reflection (van Fraassen 1984), or Jeffrey conditioning (Armendt 1980; Skyrms 1987b). In the case of reflection, propositions described by one's

11. This subsumes van Fraassen's original formulation of the reflection principle, which basically states that $\mathbb{P}[A|Y = r] = r$ for almost every ω, such that $Y(\omega) = r$ if the event $\{Y = r\}$ has positive probability (van Fraassen 1984). Our formulation of R3 goes significantly beyond this. Even if every event $\{Y = r\}$ has probability 0, it will be the case that $\mathbb{P}[A|Y = r] = r$ for almost every ω, such that $Y(\omega) = r$.

12. This is a mechanism that ideally ensures that the agent announces her fair-betting odds. It is similar to the well-known mechanism for achieving a fair division of a cake between two persons. One person cuts the cake, and the other one chooses a piece.

future degrees of belief Y for A can be used in the dynamic coherence argument only if $\sigma(Y)$ is a subset of $\mathcal{F}$. If this holds, then principle R3 or an appropriate generalization follows from dynamic coherence.

One feature of a dynamic Dutch book should be emphasized. It requires you to announce your future betting odds today; if you only announce your fair-betting odds at different times, there is no Dutch book to be had (Hacking 1967). This leads to two legitimate interpretations of future degrees of belief in this context; not surprisingly, they are the same as in the previous section. We have to view future degrees of beliefs (i.e., future betting odds) either as a plan to update one's degree of belief or as the degree of belief for A that one believes, with probability 1, is going to be one's future degree of belief.

The problem with this line of reasoning is that there exist many putative counterexamples to reflection.[13] Many of these counterexamples are aimed at reflection. But since dynamic coherence implies reflection, they are also used to argue that there must be something wrong with dynamic coherence arguments. I consider dynamic coherence first, before returning to the implications for reflection.

The structure of many counterexamples is that you expect to be in some kind of pathological situation in which your degrees of belief will be bonkers. It is then argued that you certainly do not want your degrees of belief to obey reflection, for this would require your current (sane) beliefs to conform with your future (insane) beliefs. The story of Ulysses and the Sirens is often used as an illustration. Ulysses's current beliefs are not a mixture of his anticipated future beliefs, for these will be influenced by the Sirens' song. His beliefs do not observe reflection, but it is not because his current beliefs should be any different.

Not all counterexamples have this structure. But all derive their force from using dynamic coherence arguments in a certain way. It is supposed that a rational agent is required to respond to the existence of a Dutch book by adjusting her prior so that it coheres with her posterior. This, of course, leads to absurdities. But is this a reasonable use of Dutch book arguments?

There is an interpretation of Dutch book arguments that does not lead to the conclusions drawn in the counterexamples to reflection. Several supporters of Dutch book arguments have a fairly modest purpose in mind: dynamic incoherence is used to detect inconsistencies—nothing more, nothing less. This modest purpose goes back to Ramsey (1931) and is championed, in one form or another, by various authors (e.g., Skyrms 1987a; Armendt 1993; Howson and Urbach 1993; Christensen 1996). Ramsey maintains that the existence of a Dutch book indicates inconsistencies among degrees of beliefs: "If anyone's mental condition violated these laws, his choice would

13. See Briggs (2009) for a classification and discussion of counterexamples.

depend on the precise form in which the options were offered him, which would be absurd. He could have a book made against him by a cunning bettor and would then stand to lose in any event" (1931, 182). This famous quote suggests that being taken advantage of is not the philosophically important point about a Dutch book, however unfortunate it might be; it is rather that you assign different betting odds (degrees of belief) to equivalent bets (equivalent events; see Skyrms 1987a). Dutch book arguments indicate such inconsistencies. They are tools for diagnosis.

Ramsey refers to synchronic coherence. Something similar can be spelled out for belief change—be it in terms of conditionalization, Jeffrey conditioning, or black-box learning (see Skyrms 1987b). If an agent is diachronically incoherent, then the agent has distinct fair-betting odds (degrees of belief) for equivalent bets (equivalent events). The bets are equivalent, given the underlying Boolean logic and how the agent updates beliefs—equivalent by her own standards, that is to say. Assuming the background logic and synchronic coherence, the epistemic defect brought out by the Dutch book is in this case attributable to how the agent updates beliefs.

The important question now is what a dynamically incoherent agent should do. If we take her to be only concerned about the pragmatic aspect of the Dutch book, it is reasonable to conclude that she will change her beliefs in order not to be exploitable. But if we focus on the epistemic defect that is indicated by dynamic incoherence, we are not forced to this conclusion. Instead, a dynamic Dutch book informs the agent that her degrees of beliefs are inconsistent but says nothing about how the agent should respond to this inconsistency. Prima facie, consistent degrees of belief are an epistemic virtue. But this does not imply that an agent has to establish consistency at any cost; the Dutch book does not override all other considerations that might be important for the agent.

The Ramseyan point of view leads to a tempered understanding of dynamic incoherence. A Dutch book indicates a particular epistemic defect, but it does not say anything about whether or how the agent should change her degrees of belief. The agent's response will often depend on other considerations.

Understanding dynamic coherence in this way has some important consequences for the counterexamples to reflection. In some of them, the agent might well believe that her belief change will lead her to adopt irrational beliefs. Such is the case in the example of Ulysses and the Sirens. Briggs (2009) points out that in cases like this one, an agent exhibits self-doubt. A rational agent should be able to have such beliefs, without being forced to obey principles such as reflection. But notice that this belief of the agent's—that updating her beliefs will be irrational—is correctly indicated by dynamic incoherence. The Dutch book argument is actually doing its job, and no absurd consequences can be derived since the agent is not required to avoid the Dutch book by changing her prior.[14]

Other counterexamples to reflection involve memory loss. These examples show, correctly I think, that dynamic coherence and reflection are insufficient for all-things-considered rationality. But this conclusion misses the point. The claim is not that dynamic coherence and reflection are sufficient for all-things-considered rationality. The claim is that dynamic incoherence and violations of reflection are indicators of *epistemic irrationality*. Forgetting is one kind of epistemic irrationality. But it is perfectly rational (in the all-things-considered sense) to prefer a situation in which one is slightly epistemically irrational to a situation in which one is perfectly epistemically rational but has to pay all sorts of nonepistemic costs.

There is another class of counterexamples that I will not discuss. They are similar to the case of Sleeping Beauty in that an agent can learn events that she does not already grasp. This violates our basic assumptions and leads outside the standard mathematical theory of conditional expectation. Therefore, such examples deserve a special treatment.[15]

In brief, the following view emerges from these considerations. Dynamic incoherence—understood in a tempered sense and applied to situations that fall within the scope of the theory of conditional expectation—as well as reflection are diagnostic of epistemic irrationality. The epistemic irrationality applies to how an agent updates beliefs since we have assumed that the agent is synchronically rational. Thus, as long as one ignores larger considerations, an agent cannot violate reflection and at the same time think that she will form her future degrees of belief in an epistemically rational way. If she does consider herself to be epistemically rational, then her probability measure should observe reflection.

Phrased in terms of plans, this is essentially the conclusion reached in van Fraassen (1995) as to what a violation of reflection amounts to: "the person holding this opinion cannot regard herself as following a rational policy for opinion change" (17). Similarly, Skyrms (1990) takes reflection to indicate cases of genuine learning. What my discussion adds to van Fraassen's and Skyrms's arguments is that these conclusions derive very naturally from a view of Dutch book arguments as tools for diagnosis. In addition, it allows us to largely deflate the intuitive plausibility of some counterexamples to dynamic coherence and reflection, by observing that a diagnosis of dynamic incoherence does not say anything about how an agent should respond to

14. Jeffrey (1988) also ties violations of reflection to cases in which the agent expects her belief change to be unreasonable.

15. See, e.g., Schervish, Seidenfeld, and Kadane (2004) for a thoughtful discussion of this issue.

incoherence. In fact, violations of reflection are indicative of an agent who cannot consider herself to be epistemically rational.

Much of my argument depends on the connection between fair-betting odds and degrees of belief. This connection implies that having distinct fair odds for equivalent bets is the same as having distinct degrees of belief for equivalent events. If this link does not hold, then the agent's evaluations of bets are inconsistent. While this is certainly some kind of defect, it need not count as epistemically defective (see Joyce 1998).

I think that one can make sense of the connection between betting odds and degrees of belief, in terms of measuring strength of belief by betting behavior. However, this is primarily a measurement-theoretic question that would lead us too far afield. Instead of pursuing this topic further, I introduce another approach for justifying reflection principles in which we start with numerical beliefs as a primitive concept. This approach might be more appealing to those philosophers who entertain fundamental doubts about Dutch book arguments.

4. Reflection and Expected Inaccuracy. We have seen that reflection is a feature of consistent degrees of belief, if degrees of belief are measured by fair-betting behavior. In the next two sections, we see that it is also a natural aspect of degrees of belief within the context of approximating the truth, understood in terms of minimizing expected inaccuracy. We start by considering the case of quadratic inaccuracy measures, which play an important role in accuracy-based justifications of probabilistic concepts (Greaves and Wallace 2006; Leitgeb and Pettigrew 2010a). In the following section, we see how this approach can be generalized to a natural class of inaccuracy measures.

Before I explain the details, let me try to describe the basic idea in a nontechnical way. The framework of this section is a geometrical one. We start with a space whose points are random variables; in particular, truth values of propositions and their estimates are points in this space. We can measure how much the estimate of the truth of a proposition differs from its truth value at a world, by taking the square of the difference between the value of the estimate and the truth value at this world. Taking the expectation of the squared difference relative to an agent's current probability measure yields a measure of the expected inaccuracy of the truth estimate. Now, this expectation can be used to define a distance between points in the space of random variables. With respect to this distance, it can be demonstrated that conditional probability is the best estimate of the truth value of a proposition in the sense of having minimal distance to its truth value. Most important, we shall see that this implies the martingale property $\mathbb{P}[A|Y] = Y$. In other words, the reflection principle R3 is a consequence of this approach.

It is well known that minimizing expected inaccuracy can be used to justify conditionalization (Greaves and Wallace 2006; Leitgeb and Pettigrew 2010b; Easwaran 2013). Moreover, it can be shown that, in a certain sense, conditionalization entails reflection. Hence, minimizing expected inaccuracy leads to reflection for conditionalization, as is noted by Easwaran (2013). While this is true, the arguments in this section are not merely old wine in new skins. As we have seen in section 2, the claim that conditionalization entails reflection is basically the same as the fact that an event's prior probability is equal to the expectation of its conditional probability (see eq. [4]). Reflection principles go much beyond that, however. They apply not just to conditional probabilities but to anticipated future degrees of belief in general. The fact that conditionalization entails reflection does not, by itself, allow us to conclude that reflection principles also hold for anticipated future degrees of belief that are not given by conditionalization. This requires a separate treatment.

To make the argument more precise, we start with some technicalities. Let A be an event in $\mathcal{F}$, and suppose that Y is the agent's estimate of the truth value I_A of A. Quadratic inaccuracy measures give the inaccuracy of Y as an estimate of I_A at world ω as $(I_A(\omega) - Y(\omega))^2$ up to multiplication by a positive constant. For our purposes, the constant can be ignored. The expected inaccuracy of Y as an estimate of I_A is then given by

$$E[(I_A - Y)^2]. \tag{5}$$

The expectation is taken with respect to $\mathbb{P}$.

Leitgeb and Pettigrew (2010a) provide axioms that single out quadratic inaccuracy measures as the uniquely legitimate ones. But their arguments only apply to finite probability spaces. Our use of the quadratic inaccuracy measure can therefore be justified in cases in which Ω is finite, by appealing to Leitgeb and Pettigrew's axiomatic treatment.[16] For infinite probability spaces, the results in this section should be viewed as expository. I provide a fuller treatment in the next one.

Quadratic inaccuracy measures give rise to a geometric structure that is very well known in probability theory. To introduce this structure, consider the set of all square-integrable, $\mathcal{F}$-measurable random variables (X is square integrable if $E[X^2] < \infty$). Square integrability is appropriate in the present context, for otherwise the expected inaccuracy (5) could be infinite or undefined.

16. We assume, unlike Leitgeb and Pettigrew, that $\mathbb{P}$ is a probability measure. Thus, we could use the set of axioms provided by Selten (1998) as an alternative justification of quadratic inaccuracy measures.

The set of all square-integrable random variables is a vector space over the real numbers.[17] This vector space is called $\mathfrak{L}^2$. The expectation of the product of two random variables $E[XY]$ defines an inner product on $\mathfrak{L}^2$. Variables X and Y are *orthogonal* if $E[XY] = 0$.[18] The norm associated with the inner product is $\| X \| = (E[X^2])^{1/2}$. The norm of X can be understood as the size of X. It also gives rise to a notion of distance between X and Y, by letting their distance be equal to $\| X - Y \| = E[(X - Y)^2]^{1/2}$.

The vector space $\mathfrak{L}^2$ is complete; that is, there are no "holes" in $\mathfrak{L}^2$. If we have a sequence of square-integrable random variables such that the distance between all but finitely many of them becomes arbitrarily small, the sequence converges to an element of $\mathcal{L}^2$. (More technically, with respect to the norm $\| \cdot \|$, all Cauchy sequences of elements of $\mathfrak{L}^2$ converge to an element in $\mathfrak{L}^2$.) The limit is not unique, but almost surely unique. If Y is a limit of the sequence, then any Y' with $\| Y - Y' \| = 0$ is also a limit, which means that Y and Y' are the same, except on a set of probability 0.

The geometry of this vector space is very similar to the geometry of Euclidean vector spaces, such as $\mathbb{R}^3$, with the usual dot product between vectors. Euclidean spaces can be generalized by considering spaces whose points are not real valued vectors and by using inner products other than the dot product. Such spaces are known as *Hilbert spaces*. Space $\mathfrak{L}^2$ can be viewed as a Hilbert space up to random variables that agree almost surely (Williams 1991, 65).

Suppose that $\mathfrak{K}$ is a complete vector subspace of $\mathfrak{L}^2$. Due to the Hilbert space structure of $\mathfrak{L}^2$, for any X in $\mathfrak{L}^2$, there exists a Y in $\mathfrak{K}$ that minimizes $\| X - Y \|$ such that Y is almost surely unique. Any Y' that agrees with Y except on a set of probability 0 will also minimize $\| X - Y \|$. The random variable Y is the *orthogonal projection* of X on $\mathfrak{K}$. That is, $X - Y$ is orthogonal to all random variables Z in $\mathfrak{K}$ (Williams 1991, 67).

If $\mathcal{G}$ is a sub-σ-algebra of $\mathcal{F}$, then the set $\mathfrak{G}$ of square-integrable $\mathcal{G}$-measurable random variables is a complete vector subspace of $\mathfrak{L}^2$. Whenever X is square integrable, this means that $E[X|\mathcal{G}]$ is in $\mathfrak{G}$.[19] It can be shown that $E[X|\mathcal{G}]$ is the orthogonal projection of X on $\mathfrak{G}$ (Williams 1991, 85). This implies that $E[X|\mathcal{G}]$ is the closest random variable to X among all $\mathcal{G}$-measurable random variables; see figure 1 for an illustration.

17. If X and Y are square integrable, then so is $\lambda X + \mu Y$ for all real numbers λ, μ.

18. The inner product can be used to define a generalized notion of "angle" between two random variables in $\mathfrak{L}^2$. Just as the dot product of two vectors in $\mathbb{R}^3$ is the product of their norms times the cosine of the angle between them, one can think of the inner product of two random variables as the product of their norms times the cosine of the "angle" between them; this cosine turns out to be the correlation between the variables.

19. If X is square integrable, then so is the conditional expectation $E[X|\mathcal{G}]$.

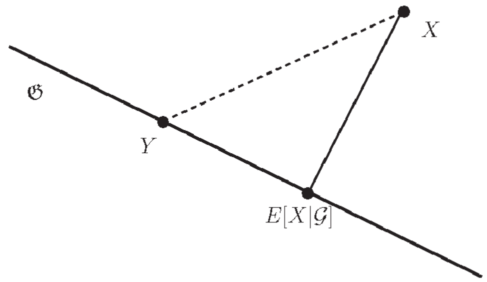

Figure 1. Set $\mathfrak{G}$ is the complete subspace of $\mathcal{G}$-measurable square-integrable random variables, and $E[X|\mathcal{G}]$ is the orthogonal projection of X on $\mathcal{G}$. If Y is not the orthogonal projection of A on $\mathfrak{G}$, then it does not minimize the distance to X. The set $\mathcal{G}$ can be equal to the set $\mathfrak{Y}$ which contains all $\sigma(Y)$-measurable random variable.

Closeness refers to the norm $\|\cdot\|$. By definition, then, $E[(X - E[X|\mathcal{G}])^2]$ minimizes the expected square error $E[(X - Y)^2]$ for all $\mathcal{G}$-measurable Y, and $E[X|\mathcal{G}]$ is the almost unique minimum. The expected value

$$E[(X - Y)^2] \tag{6}$$

is a generalization of the expected inaccuracy in (5), where X is the true value of some quantity and Y is the agent's estimate of X after learning which event in $\mathcal{G}$ obtained. Thus, $(X(\omega) - Y(\omega))^2$ is the inaccuracy of her estimate at world ω. In this case, the agent minimizes expected inaccuracy by choosing Y to be almost equal to $E[X|\mathcal{G}]$.

Leitgeb and Pettigrew prove that conditional probability minimizes expected inaccuracy in finite probability spaces (2010b, theorem 3). The considerations above yield a version of this theorem for general countably additive probability spaces. In a first step, observe that the argument in the previous paragraph implies that the conditional probability $\mathbb{P}[A|\mathcal{G}]$ is the $\mathcal{G}$-measurable random variable that minimizes the expected inaccuracy $E[(I_A - Y)^2]$. However, it is not in general possible to choose $\mathbb{P}[\cdot|\mathcal{G}]$ so that it is a countably additive probability measure almost surely. To get a full analogue of Leitgeb and Pettigrew's theorem, $\mathbb{P}[\cdot|\mathcal{G}]$ also needs to be a *regular conditional probability*: $\mathbb{P}[\cdot|\mathcal{G}]$ is a regular conditional probability if it is a probability measure on $\mathcal{F}$ for almost every ω and $\mathbb{P}[A|\mathcal{G}]$ is $\mathcal{F}$-measurable for each A in $\mathcal{F}$.

It is well known that $\mathbb{P}[\cdot|\mathcal{G}]$ is a regular conditional probability under the fairly natural assumption that Ω is a Polish space, that is, a separable completely metrizable topological space. Together with the smallest σ-algebra that contains all open sets of Ω, such spaces are known as *standard measure spaces*. Any finite Ω is Polish, as is any Ω that has the same structure as a real metric vector space. In general, to say that Ω is a metrizable topological space and that $\mathcal{F}$ is the smallest σ-algebra containing all open sets of Ω is just to say that it is possible to give a metric distance function between points in Ω such that $\mathcal{F}$ contains every ball of any radius around every point in Ω. One can think of this in terms of a distance function between possible worlds. The distance function need not be unique. All that is required is that there exists one such distance function. To say that Ω is completely metrizable is to say Ω is complete with respect to the distance function; there are no "holes" in Ω (just like there are no "holes" in $\mathfrak{L}^2$). Separability requires that any two points in a topological space are contained in open sets that do not intersect. This could plausibly fail in relevant probability spaces. In sum, though, Polish spaces are the spaces most often encountered in applications of probability theory. Thus, for most probabilistic applications $\mathbb{P}[\cdot|\mathcal{G}]$ will almost surely be a probability measure on $\mathcal{F}$.

The same kind of reasoning as in the case of probabilities conditional on $\mathcal{G}$ can be used to show that minimizing expected inaccuracy (6) entails reflection. Let the $\mathcal{F}$-measurable random variable X again be the true value of some quantity. Because $E[X]$ minimizes expected inaccuracy (6) with respect to the trivial σ-algebra $\{\emptyset, \Omega\}$ one's current best estimate of X is its expectation.

Let Y be your anticipated new estimate of X after a black-box learning experience. As before, we assume that Ω is rich enough for Y to be well-defined and that $\sigma(Y)$ is a sub-σ-algebra of $\mathcal{F}$. In this case $\sigma(Y)$ gives rise to the vector subspace $\mathfrak{Y}$ of $\mathfrak{L}^2$ of all $\sigma(Y)$-measurable random variables.

Now suppose that

$$E[X|Y] \neq Y$$

on a set of positive probability. Then Y does not minimize expected inaccuracy (6) among all random variables in $\mathfrak{Y}$. (See fig. 1 for an illustration.) It is important to observe that all random variables in $\mathfrak{Y}$ are available to the agent, in the sense of availability used by Greaves and Wallace (2006) and Easwaran (2013). The random variables that are available to the agent are those that do not depend on any conceptual resources beyond those implicit in Y. These conceptual resources are given by $\sigma(Y)$. Hence, if Y is available as an update of one's estimate of X, so should any other random variable in $\mathfrak{Y}$, whereas it need not be the case that the agent is able to express future degrees of belief that go beyond $\sigma(Y)$. Consider another random variable Y', and suppose, for example, that there is a set B in $\sigma(Y')$ that is not a member of $\sigma(Y)$. Then the conceptual resources implicit in Y' go beyond those of Y, and it is not the case that Y' is available for the agent whenever Y is.

Among all random variables in $\mathfrak{Y}$, $E[X|Y]$ minimizes expected inaccuracy. If $Y \neq E[X|Y]$ is the agent's estimate of X, then she does not approximate the true values X as closely as she could given the available random variables. If, however,

$$E[X|Y] = Y \tag{7}$$

almost surely, then expected inaccuracy is minimized. Equation (7) is the martingale property and, hence, a general version of principle R3. We also get versions of principles R1 and R2 since $E[Y] = E[E[X|Y]] = E[X]$ (the last equality is a standard property of conditional expectations). It follows that minimizing expected inaccuracy in terms of (6) entails reflection for random variables.

Reflection principles for degrees of belief can be obtained by the familiar substitution $X = I_A$. The foregoing arguments then imply that Y minimizes expected inaccuracy if and only if

$$\mathbb{P}[A|Y] = Y \tag{8}$$

almost surely. Hence, by minimizing expected inaccuracy you update your degrees of belief as if you would condition on Y (cf. Good 1981). It also follows that $E[Y] = E[\mathbb{P}[A|Y]] = \mathbb{P}[A]$. Taken together, this again leads to versions of the three reflection principles R1–R3.

5. Generalizations. In the previous section we assumed that inaccuracy measures are quadratic. One might wonder to what extent the results for reflection depend on this assumption. This question is especially important in the absence of a more principled justification of quadratic inaccuracy measures for general probability spaces. Some theorems that are proven in Banerjee, Guo, and Wang (2005) can be used to show that the results of the previous section continue to hold for a large class of inaccuracy measures called *Bregman distance functions*, which can be derived from a salient condition on estimates, namely, a generalization of the condition of *propriety* or *immodesty* (see conditions [10] and [11] below). Hence, reflection principles and, in particular, the martingale property follow from a quite natural epistemic condition on the accuracy of estimates.

Suppose that $\phi : D \to \mathbb{R}$ is a strictly convex differentiable function, where D is an interval in $\mathbb{R}$. Then the Bregman distance function $B_\phi : D \times D \to \mathbb{R}$ is defined as

$$B_\phi(x, y) = \phi(x) - \phi(y) - (x - y)\phi'(y),$$

where ϕ' denotes the derivative of ϕ. The function B_ϕ is the difference between the value of ϕ at x and the value of the first-order term of the Taylor expansion of ϕ around y evaluated at x. The function $\phi(x) = x^2$, for example, gives rise to the distance function $(x - y)^2$. Thus, the quadratic inaccuracy measure is a Bregman distance function. Other examples include the Kullback-Leibler divergence or the Itakura-Saito distance (Banerjee et al. 2005).

Suppose that X is a $\mathcal{F}$-measurable random variable for which both $E[X]$ and $E[\phi(X)]$ are finite and that B_ϕ is a Bregman distance function. The expected inaccuracy of the estimate Y is now given by

$$E[B_\phi(X, Y)]. \tag{9}$$

It can be shown that $E[X|\mathcal{G}]$ is the (almost surely) unique minimizer of $E[B_\phi(X, Y)]$ among all Y that are $\mathcal{G}$-measurable (see Banerjee et al. [2005] for a proof).

This result allows us to draw the same conclusions for all Bregman distance functions as for quadratic inaccuracy measures. In particular, if Y is the agent's future estimate of X that minimizes expected inaccuracy in terms of (9), then $E[X|Y] = Y$ almost surely. From this the reflection principles for conditional expectations and conditional probabilities follow.

One can also show that Bregman distance functions are, under rather mild regularity conditions, the only ones for which the conditional expectation $E[X|\mathcal{G}]$ is closest to X among all $\mathcal{G}$-measurable random variables (Banerjee et al. 2005). Let $F: D \times D \rightarrow \mathbb{R}$ be a nonnegative continuous function such that F is continuously differentiable in its first argument and $F(x, x) = 0$ for all $x \in D$. If $E[X|\mathcal{G}]$ minimizes $E[F(X, Y)]$ among all Y that are $\mathcal{G}$-measurable, then F is a Bregman distance function for some strictly convex function ϕ.

Why should one choose a Bregman distance function as one's measure of inaccuracy? The last result does not provide a useful rationale, for it assumes our desired conclusion: that a random variable's conditional expectation is its most accurate estimate. However, Banerjee et al. (2005) prove something stronger than this. Suppose that for all measurable random variables X and for all constant random variables Z, such that $E[X] \neq Z$,

$$E[F(X, E[X])] < E[F(X, Z)]. \tag{10}$$

That is, $E[X]$ is the unique minimizer of $E[F(X, Z)]$ among all constant random variables Z. If $X = I_A$, this means that for any $A \in \mathcal{F}$ and for any constant random variable Z such that $\mathbb{P}[A] \neq Z$,

$$E[F(I_A, \mathbb{P}[A])] < E[F(X, Z)]. \tag{11}$$

Therefore, measured in terms of F, your degree of belief $\mathbb{P}[A]$ is the most accurate constant estimate of I_A in expectation. Banerjee et al. (2005) prove that if (10) holds for a continuous nonnegative function F with $F(x, x) = 0$, and if F is continuously differentiable in its first argument, then F is a Bregman distance function for some strictly convex function ϕ. Hence, if (10) holds, then $E[X|Y]$ is the closest random variable to X among all $\sigma(Y)$-measurable random variables with respect to F. We see that, once again, our reflection principles follow, this time from condition (10).

Conditions such as (10) and (11) are known as propriety conditions and have been much discussed in the literature (Greaves and Wallace 2006; Joyce 2009; Easwaran 2013). Degrees of beliefs obeying a condition like (11) are often called *immodest*. From the point of view of measuring inaccuracy, modesty appears to be a vice. Modest degrees of beliefs are self-undermining; they recommend degrees of beliefs other than the ones the agent is currently holding as being superior. Put differently, the choice of a function F that violates (11) is inconsistent with regarding all one's current truth estimates as maximally accurate. The same can be said about constant estimates of random variables in (10). Modest estimates violate (10) and are therefore self-undermining. If $E[X]$ is not the best constant estimate of X, then adopting a different estimate would be superior in the light of the agent's own probabilities.

Constant estimates are significant in that they are always available to the agent (Greaves and Wallace 2006). Only constant estimates are measurable relative to the trivial σ-algebra $\{\varnothing, \Omega\}$. One therefore does not need any information about the structure of the probability space in order to form a constant estimate.

The results of this and the previous section bear some resemblance to the justifications of conditionalization in Greaves and Wallace (2006) and Easwaran (2013). In both of these articles, conditions very similar to (11) play an important role in showing that conditional probabilities minimize expected inaccuracy when planning to update on the outcomes of an experiment. An experiment is taken to be a partition of Ω. Our new results show that general conditional expectations also minimize inaccuracy. That is, first, the outcomes need not constitute a partition of Ω but can be a σ-algebra, which is a more general structure. And, second, this σ-algebra can be generated by the future estimate of a quantity. Easwaran has pointed out that reflection in the sense of R2 holds for conditional probabilities given an experiment. Within our framework, we obtain more general reflection principles—ones that also apply to general future estimates of X.

Should (10) be viewed as a rationality requirement? This question boils down to asking whether a rational agent can hold constant estimates that are self-undermining. If we consider an agent whose only goal is to approximate the truth, then this should clearly not be the case. Such an agent would, on learning that her constant estimates are not the most accurate ones, either change her constant estimates or have doubts about the appropriateness of the function F as a measure of inaccuracy. In the framework of expected accuracy, it is natural that degrees of belief are estimates of truth values and, more generally, that expectations are estimates of random variables. Violations of (10) and (11) thus amount to saying that the distance measure F itself precludes certain values from being one's current estimates.

This is not to deny that (10) is a demanding requirement. It would be easy to find examples in which it does not seem to be reasonable for an agent to meet (10), regardless of the costs involved. Like dynamic coherence, maximally accurate constant estimates in the sense of (10) should be viewed as a prima facie requirement of rationality. Propriety does not override all considerations in a larger all-things-considered context. But within the confines of epistemic rationality, we can again conclude that reflection is a rationality requirement, if truth approximation is taken to be the standard of rationality.

6. Concluding Remarks. In section 2, we saw that reflection principles are closely tied to the general concept of expectation conditional on a σ-algebra. Importantly, this does not depend on where the σ-algebra comes from. The σ-algebra might describe the outcomes of an experiment. But it can also be generated by your future beliefs—which is an experiment of sorts in the context of black-box learning.

From this we were led to the conclusion that reflection is a basic feature of rational opinion change. Two epistemic principles that require this are dynamic coherence and accuracy. These two approaches highlight different epistemic virtues of reflection. Dynamic coherence brings out consistency, whereas accuracy emphasizes approximating the truth. In many situations, other considerations also influence one's anticipated future estimates. But if one only cares about consistency or approximating the truth, then reflection is a requirement for rational change of opinions.

REFERENCES

Adams, Ernest W. 1962. "On Rational Betting Systems." *Archiv für mathematische Logik und Grundlagenforschung* 6:7–29, 112–28.
Armendt, Brad. 1980. "Is There a Dutch Book Argument for Probability Kinematics?" *Philosophy of Science* 47:583–88.
———. 1993. "Dutch Books, Additivity, and Utility Theory." *Philosophical Topics* 21:1–20.
Arntzenius, Frank. 2003. "Some Problems for Conditionalization and Reflection." *Journal of Philosophy* 100:356–70.
Bacchus, Fahiem, Henry E. Kyburg, and Mariam Thalos. 1995. "Against Conditionalization." *Synthese* 85:475–506.
Banerjee, Arindam, Xin Guo, and Hui Wang. 2005. "On the Optimality of Conditional Expectation as a Bregman Predictor." *IEEE Transactions on Information Theory* 51:2664–69.
Bovens, Luc. 1995. "'P and I Will Believe That Not-P': Diachronic Constraints on Rational Belief." *Mind* 104:737–60.
Briggs, Rachael. 2009. "Distorted Reflection." *Philosophical Review* 118:59–85.
Christensen, David. 1991. "Clever Bookies and Coherent Beliefs." *Philosophical Review* 100: 229–47.
———. 1996. "Dutch Books Depragmatized: Epistemic Consistency for Partial Believers." *Journal of Philosophy* 93:450–79.

Easwaran, Kenny. 2013. “Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection.” *Philosophy of Science* 80:119–42.

Goldstein, Michael. 1983. “The Prevision of a Prevision.” *Journal of the American Statistical Association* 78:817–19.

Good, Irving J. 1981. “The Weight of Evidence Provided by an Uncertain Testimony or an Uncertain Evidence.” *Journal of Statistical Computation and Simulation* 13:56–60.

Greaves, Hilary, and David Wallace. 2006. “Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility.” *Mind* 115:607–32.

Hacking, Ian. 1967. “Slightly More Realistic Personal Probability.” *Philosophy of Science* 34:311–25.

Howson, Colin, and Peter Urbach. 1993. *Scientific Reasoning: The Bayesian Approach*. 2nd ed. La Salle, IL: Open Court.

Jeffrey, Richard C. 1988. “Conditioning, Kinematics, and Exchangeability.” In *Causation, Chance, and Credence*, vol. 1, ed. Brian Skyrms and William L. Harper, 221–55. Dordrecht: Kluwer.

Joyce, James M. 1998. “A Nonpragmatic Vindication of Probabilism.” *Philosophy of Science* 65: 575–603.

———. 2009. “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief.” In *Degrees of Belief*, ed. Franz Huber and Christoph Schmidt-Petri, 263–97. New York: Springer.

Leitgeb, Hannes, and Richard Pettigrew. 2010a. “An Objective Justification of Bayesianism.” Pt. 1, “Measuring Inaccuracy.” *Philosophy of Science* 77:201–35.

———. 2010b. “An Objective Justification of Bayesianism.” Pt. 2, “The Consequences of Minimizing Inaccuracy.” *Philosophy of Science* 77:236–72.

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

Maher, Patrick. 1992. “Diachronic Rationality.” *Philosophy of Science* 59:120–41.

Ramsey, Frank P. 1931. “Truth and Probability.” In *Foundations of Mathematics and Other Essays*, ed. Richard B. Braithwaite. New York: Harcourt Brace. Repr. in *Studies in Subjective Probability*, ed. Henry E. Kyburg and Howard E. Smokler (Huntington, NY: Krieger, 1964).

Schervish, Mark J., Teddy Seidenfeld, and Joseph B. Kadane. 2004. “Stopping to Reflect.” *Journal of Philosophy* 101:315–22.

Selten, Reinhard. 1998. “Axiomatic Characterization of the Quadratic Scoring Rule.” *Experimental Economics* 1:43–61.

Skyrms, Brian. 1987a. “Coherence.” In *Scientific Inquiry in Philosophical Perspective*, ed. Nicholas Rescher, 225–342. Pittsburgh: University Press of America.

———. 1987b. “Dynamic Coherence and Probability Kinematics.” *Philosophy of Science* 54:1–20.

———. 1990. *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.

———. 2006. “Diachronic Coherence and Radical Probabilism.” *Philosophy of Science* 73:959–68.

Talbott, William. 1991. “Two Principles of Bayesian Epistemology.” *Philosophical Studies* 62: 135–50.

Teller, Paul. 1973. “Conditionalization and Observation.” *Synthese* 26:218–58.

van Fraassen, Bas C. 1984. “Belief and the Will.” *Journal of Philosophy* 81:235–56.

———. 1995. “Belief and the Problem of Ulysses and the Sirens.” *Philosophical Studies* 77:7–37.

Weisberg, Jonathan. 2007. “Conditionalization, Reflection, and Self-Knowledge.” *Philosophical Studies* 135:179–97.

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