

Bayesianism I: Introduction and Arguments in Favor

Kenny Easwaran*

University of Southern California

Abstract

Bayesianism is a popular position (or perhaps, positions) in the philosophy of science, epistemology, statistics, and other related areas, which represents belief as coming in degrees, measured by a probability function. In this article, I give an overview of the unifying features of the different positions called ‘Bayesianism’, and discuss several of the arguments traditionally used to support them.

Bayesianism is a collection of positions in several related fields, centered on the interpretation of probability as something like degree of belief, as contrasted with relative frequency, or objective chance. However, Bayesianism is far from a unified movement. Bayesians are divided about the nature of the probability functions they discuss; about the normative force of this probability function for ordinary and scientific reasoning and decision making; and about what relation (if any) holds between Bayesian and non-Bayesian concepts.

The core ideas that most Bayesians can agree on are roughly the following, which I will discuss in more detail later on:

1. There is an important mental attitude of *degree of belief* or *credence* that can often be given numerical values.
2. For an agent to be perfectly rational, her degrees of belief must obey the axioms of probability theory.
3. Conditionalization (to be described later), or some close relative, is the standard way beliefs change over time.

The first claim is the central claim of Bayesianism, though there is disagreement about whether degree of belief is the *only* important doxastic and epistemic attitude, or whether others are important as well. The second claim then gives rationality conditions that degrees of belief must satisfy at a time (synchronically), while the third gives conditions on how they must relate across time (diachronically). Omitting the third claim gives a position called ‘probabilism’ that is often defended even by opponents of full Bayesianism. The fact that degrees of belief are represented mathematically allows for the precision and power of probabilism and Bayesianism, but it is also the source of much skepticism towards these views among other philosophers.

I will leave the primary discussion of the first claim for later on, because it is one of the sources of disagreement among Bayesians. But it seems clear that although I believe that the sun will rise tomorrow, and that it will not rain in Los Angeles tomorrow, I am far more confident about the sun’s rising. Thus, there must be some notion of degree of belief alongside the traditional epistemic notions of belief, justification, and knowledge, though its relation to them is unclear.

As for the second claim, there is no single formulation of the synchronic axioms degrees of belief must satisfy. The most canonical set comes from Kolmogorov (1950), but his axioms are a pure mathematical abstraction, and require interpretation to be applied to Bayesianism. Kolmogorov defines an abstract mathematical structure, which he calls a ‘field of probabilities’. This is a set E , together with a collection \mathcal{F} of subsets of E , and a function P assigning real numbers to the members of \mathcal{F} . He gives a characterization of \mathcal{F} ,¹ but the axioms that are often described as ‘Kolmogorov’s axioms’ are the three axioms on P . First, the values P takes are always non-negative. Second, $P(E) = 1$. Third, if A and B are disjoint sets (so that they have no elements in common), then $P(A \cup B) = P(A) + P(B)$, where $A \cup B$ is the set of all elements that are in A or B or both.²

Now clearly, these axioms are very distant from anything remotely mental or epistemic – they just describe an abstract mathematical set. However, if the set E is taken to be the set of epistemic possibilities for an agent, and any proposition is identified with the set of epistemic possibilities on which it is true, then we have the start of a plausible interpretation. If we further assume that a conjunction is true in any possibility in which the two conjuncts are, and that a negation is true in every possibility in which the basic proposition isn’t, then these propositional connectives correspond to the set-theoretic operations of intersection and complement, respectively. Of course, there are worries about just what epistemic possibilities might be, and whether they really exist. In order to sidestep these issues, many authors prefer to use axioms like those in Popper (1959) or Hailperin (1984) that just use syntactic properties of the objects of belief, rather than set-theoretic properties. However, I will argue later that there are some advantages to the Kolmogorov formulation, even if it seems implausible to identify propositions with sets of some sort.

There is also a further part to the synchronic characterization of degree of belief, which is the notion of *conditional* probability. This notion of conditional probability is intended to represent something like an agent’s degree of belief in B on the supposition that A is true, written $P(B | A)$. Some mathematicians take conditional probability to be *defined* by the formula $P(B | A) = \frac{P(B \& A)}{P(A)}$, but Hájek (2003) argues that conditional degree of belief is a pre-theoretic notion as much as degree of belief, and therefore we must analyze it, rather than just defining it mathematically. Since $B \& A$ entails A , we must have $P(B \& A) \leq P(A)$, and thus if $P(A) = 0$, then this formula says that $P(B | A) = 0/0$, which is undefined. To avoid this problem, we might weaken the ratio formula to just say that $P(B | A)P(A) = P(B \& A)$ (though we may need to supplement this with further conditions on $P(B | A)$ when $P(A) = 0$). This is the strategy taken by Popper (1959) and others in giving alternate accounts of the relation between conditional and unconditional probability. At any rate, conditional probability (although it is a purely synchronic notion) plays an important role in the diachronic norms of Bayesianism, as well as in Bayesian confirmation theory, and many other applications of probability theory.

An important theorem concerning conditional probability is Bayes’ Theorem, which in one form states that whenever the relevant probabilities are all non-zero,

$$P(A | B) = P(B | A) \frac{P(A)}{P(B)}.$$

Note that this theorem deals only with the agent’s beliefs and conditional beliefs at one time (i.e., it is totally synchronic), so it is *not* an update rule. Nothing about this theorem is distinctively Bayesian (after all, it doesn’t require that the probabilities discussed are

degrees of belief, rather than chances or frequencies), but as I will mention later, it does help illuminate several points in Bayesian confirmation theory. However, Chapter 1 of Earman (1992) points out that the paper by Rev. Bayes in which this theorem was first proved also marks some of the earliest thinking about degrees of belief and their rational update, so that ‘Bayesianism’ is an appropriate name for the view after all.

The third claim of Bayesianism gives the diachronic rules for rationality. The basic rule that is most endorsed is called ‘conditionalization’. This rule says that updates to credences are triggered by a basic learning event (e.g., making an observation), and that if A is the strongest proposition that is learned, then the agent’s new credence in any proposition is given by $P_{new}(B) = P_{old}(B | A)$. It is straightforward to verify that if P_{old} satisfied the Kolmogorov axioms, and if $P_{old}(A) \neq 0$, then P_{new} will also satisfy the Kolmogorov axioms and $P_{new}(A) = 1$. However, once a proposition has reached probability 1, no further update by conditionalization can change this status (unless we conditionalize on something with probability 0).

For this reason, many Bayesians instead endorse a modified form of conditionalization due to Richard Jeffrey. This account allows the agent’s degree of belief in A to increase by some amount as a result of the basic learning event, but without actually reaching 1. On this approach, $P_{new}(B) = P_{old}(B | A)P_{new}(A) + P_{old}(B | \neg A)P_{new}(\neg A)$.³ There are many interesting issues about Jeffrey updating that I can’t discuss here. For instance, Diaconis and Zabell (1982) shows how this can be seen as a version of traditional conditionalization if one adds particular new propositions to the language. Field (1978) and Wagner (2002) discuss what it might mean to receive ‘the same evidence’ in two different Jeffrey updates, and whether Jeffrey conditionalization can satisfy the thought that receiving the same evidence in two different orders should result in the same final credences.

Clearly, these are not the only ways that ordinary rational agents update their beliefs – people sometimes forget, and there are further complexities with other sorts of complicated situations. I can only mention a few here. A famous puzzle governing a strange sort of update is the Sleeping Beauty problem (Elga 2000), which raises issues about the distinction between possible worlds and epistemic possibilities, as well as the extent to which self-locating information can be relevant to non-self-locating beliefs. Some particularly interesting responses are Lewis (2001), Hitchcock (2004), Bradley and Leitgeb (2006), Meacham (2008), and Titelbaum (2008). The ‘Dr. Evil puzzle’, which also turns on the relevance of self-locating information to non-self-locating beliefs, is discussed in Elga (2004) and the response Weatherson (2005). There are also other problems, involving forgetting, loss of track of time, and the like, discussed in Arntzenius (2003) among other places.

1. Arguments for Bayesianism

1.1. DUTCH BOOK ARGUMENTS

The most prominent sort of argument in favor of Bayesianism is the ‘Dutch book’ argument. However, this argument does not support all aspects of Bayesian doctrine – it makes several assumptions about the mental states of rational agents, and then uses those assumptions to argue that rational agents must conform to the axioms of probability theory, and must update by conditionalization. The argument is generally attributed to both Ramsey (1926) and de Finetti (1937), although Ramsey only mentions the Dutch book argument as an afterthought to a representation theorem argument (about which more later).

The Dutch book argument starts by making the assumption that for every proposition (including conjunctions and negations of other propositions), there is some price that an agent considers fair for a bet that pays \$1 if the proposition is true and nothing if the proposition is false. 'Fair' in this setting means that the agent is willing to sell such a bet for any amount greater than the fair price, and she is willing to buy such a bet for any amount less than the fair price. Additionally, it is assumed that an agent's willingness to buy or sell one bet is unaffected by which other bets she has already bought or sold. Given these assumptions, it is a theorem that if these prices do not satisfy Kolmogorov's axioms, then there is a finite set of bets (called a 'Dutch book') the agent is willing to accept that guarantees the agent will lose money.⁴ Conversely, if these prices do satisfy Kolmogorov's axioms, then the agent is not susceptible to a standard Dutch book, though as pointed out in Hájek (2008), the agent might be susceptible to various generalized Dutch books.

We must be very careful in stating just what it means to 'guarantee' a loss of money – if *A* is false, then it looks like any willingness to pay for a bet on *A* will guarantee a loss of money. However, the notion of 'guarantee' here must be subjective – it must mean that if the agent countenances epistemic possibilities on which *A* is true, then *she* can't guarantee that she will lose money, and this is the sense of 'guarantee' that is needed. On this subjective notion of guarantee, it seems clear that a guaranteed loss of money is irrational, and thus this argument seems to show that a rational agent must satisfy Kolmogorov's axioms.⁵ Authors who prefer a more syntactic axiomatization of probability must also modify this notion of guarantee, but the prospects seem more problematic. This notion of guarantee is sometimes spelled out in terms of a 'Dutch bookie' who actually makes the bets to take advantage of the agent, who must not know anything more than the agent herself – but I think the purely internal arguments are more convincing as an illustration of the inconsistency of the agent's preferences.⁶

Now the assumptions of the argument are clearly unrealistic – it seems that most real agents are risk-averse (meaning that there is a gap between the price at which they're willing to buy a bet and the price at which they're willing to sell it), and most agents' willingness to buy and sell bets is strongly affected by whether they hold other bets that are complementary (after all, this is why people buy insurance). But if an agent had a uniform fair price for buying and selling, and if this price were determined solely by her degree of belief in the proposition, then these assumptions would be true, and furthermore the fair prices would provide a way to assign numerical values to degrees of belief. In fact, some followers of Ramsey and de Finetti have made the further claim that this price *is* the degree of belief, in order to give operational meaning to this theoretical term – on this interpretation, the assumptions of the Dutch book argument just are the first claim of Bayesianism together with a sort of 'package principle' requiring that an agent be willing to buy a set of bets if she is willing to buy every bet in the set.

However, Christensen (1996) points out that we don't need to make the full unrealistic assumptions for the argument to go through. First, we can just assume that *there is* a Bayesian notion of degree of belief, and that this notion *justifies* the agent in deeming bets favorable (regardless of whether she would be willing to accept them or not). Then we assume that an agent is justified in deeming a combination of bets favorable if she is justified in deeming each one individually favorable, and finally that a rational agent is *never* justified in deeming favorable a combination of bets that guarantees her loss. The Dutch book theorem shows that an agent will violate Kolmogorov's axioms iff there is a Dutch book, each of whose bets she is individually justified in deeming fair, and thus a rational agent will never violate Kolmogorov's axioms.

To justify the further claims about *conditional* probability, the argument needs to be extended. We can define a bet on A conditional on B to be one that pays \$1 if both A and B are true, pays nothing if A is false but B is true, and in which the cost of the bet is refunded if B is false. If we then assume that the agent's conditional degree of belief $P(A | B)$ specifies the fair price for this bet, then she is vulnerable to a conditional Dutch book iff she violates $P(A | B)P(B) = P(A \& B)$.

As arguments for the diachronic claims of Bayesianism, Dutch books are much more problematic. Instead of considering a set of bets that the agent is willing to accept all at once, we must consider a series of bets that an agent is willing to accept over time (with payoffs delayed until the end of course, so that the agent doesn't get any information from them) such that her final outcome is guaranteed to be a loss. But we must be very careful with such sequences of bets.

For instance, there is an apparent Dutch book argument against an agent *ever* changing her degrees of belief. If there is one time at which an agent's degree of belief in A is p , and another time at which it is $p - \epsilon$, then at the one time she is willing to pay $\$p - \epsilon/3$ (since this is less than p) for a bet that pays \$1 iff A is true, and at the other time she is willing to sell such a bet for $\$p - 2\epsilon/3$ (since this is more than $p - \epsilon$). Since the payoffs (if any) of these two bets cancel out, the net result is a loss of $\$ \epsilon/3$.

However, appealing to the notion of 'guarantee' from above can clear this up. This sequence of bets only guarantees a loss if *every* epistemic possibility that the agent countenances at the start involves this change of credence. If some possibilities involve an increase in credence and others involve a decrease in credence, then her initial willingness to bet doesn't guarantee a loss – she may end up in a state where she later sees that she is guaranteed a loss, but this is a risk that comes with any gamble. This argument does, however, illustrate an important constraint on updating known as 'reflection', one version of which states that if an agent is certain about what her future degrees of belief will be, then she should already have those degrees of belief. This constraint was introduced in van Fraassen (1984), and has been much discussed in subsequent literature (see Christensen 1991; van Fraassen 1995, 1999; Elga 2007, among others).

Dutch book arguments for update by conditionalization go back to Teller (1973) (though Teller attributes the argument itself to Lewis). Teller's argument only applies to a case where an agent knows in advance that she will learn which piece of evidence E_i is true, where it is also known in advance that exactly one of the E_i is true. If (say) she has non-zero degree of belief in E_1 , and yet also knows that on learning E_1 her degree of belief in some proposition A will be less than $P(A | E_1)$, then she is willing to buy a conditional bet on A given E_1 , together with a small bet on E_1 , and if E_1 is true then she will sell a bet on A , such that together the three bets guarantee a loss. (Armendt (1980) generalizes this argument to show that even if she doesn't learn which E_i is true, if the E_i form the partition whose credence shift drives the rest of the update, then the agent must update by Jeffrey conditionalization.) The conditions under which Teller's argument is valid are not totally general, but this is to be expected, because even most Bayesians admit that conditionalization is not the most general rule for updating degrees of belief, even though conditionalization and its generalized version by Jeffrey are generally considered the two most important procedures.

Thus, Dutch book arguments have been used to support the synchronic and diachronic claims of Bayesianism, in addition to helping give meaning to the primary claim that belief can come in degrees. However, the assumptions about fair betting prices that are needed seem unduly strong to many people. Additionally, the inconsistency they point out is an inconsistency of preferences, and not of beliefs, so many are worried that this

can't be an argument for an epistemic conclusion.⁷ Diachronic Dutch books also raise several new worries (Christensen 1991), as do Dutch books involving infinitely many bets (as in McGee 1999 and Williamson 1999). Thus, a lot of skepticism surrounds Dutch book arguments, so other arguments for Bayesianism are often considered as well.

1.2. REPRESENTATION THEOREMS

Another sort of argument that has been quite prominent in favor of Bayesianism is an argument by a 'representation theorem'. Probably the best reconstruction of the primary argument in Ramsey (1926) is as a representation theorem argument, and versions of it are stated more fully in Savage (1954), Jeffrey (1965), and especially Joyce (1999). Rather than assuming that agents have fair betting prices, one just makes assumptions about an agent's preferences among various actions. However, these preferences must cover *all* actions and not just those traditionally considered 'bets'.

Some of these assumptions are primarily structural, such as the assumption that agents have preferences among all 'actions' that can be defined by a function from epistemic possibilities to 'outcomes', regardless of whether or not such 'actions' are really available for an agent. This seems reasonable for straightforward bets – if I can place a bet where I win if p is true and I lose if it is false, then I can surely place the converse bet. But since we want to discuss *all* of an agent's preferences, this means that if there's an action (say, leaving your umbrella at home) that gives you a soggy walk if it rains, and a dry unencumbered walk otherwise, then there must also be some 'action' that gives you a dry unencumbered walk if it rains and a soggy walk if it doesn't. While this seems like quite an unrealistic 'action', the requirement is just that the agent have some sort of hypothetical preference ranking of this action compared to all other (actual or hypothetical) actions.

The other assumptions are more substantively about preferences. A few common assumptions are transitivity (if an agent prefers action A to action B , and B to C , then she must prefer A to C) and dominance (if an agent prefers the action that guarantees outcome O_1 to the action that guarantees the outcome O_2 , then she must prefer A' to A , where A' differs from A only in causing outcome O_1 in some situations in which A leads to outcome O_2).

For each relevant set of assumptions, one can prove a theorem that any agent whose preferences obey these axioms must be representable in a special way. In particular, there is a function U from the set of outcomes to the real numbers, and a function P from the propositions to the real numbers in $[0,1]$, such that the agent prefers action A to action B iff $\sum xP(U(A) = x) > \sum xP(U(B) = x)$, where ' $U(A) = x$ ' is the proposition stating that action A gives rise to an outcome O such that $U(O) = x$.⁸ If the assumptions made are strong enough, then one can show that in fact the function P is unique and satisfies the Kolmogorov axioms, and there is only one such function U up to positive scalar multiplication and an additive constant. Then the theorist makes one final leap and says that U represents the agent's 'utilities' (how desirable she finds various outcomes) and P represents her probabilities (how believable she finds various propositions). This identification is supported by the intuitive fact that preferences govern actions just as desires and beliefs do, and that the calculation $\sum xP(U(A) = x)$ is the standard 'expected utility' used in decision theory to represent the notion of rational preference between actions.

There are several points at which such arguments have been criticized. Zynda (2000) has shown that alternate representations are also compatible with the axioms on preference, raising the question of why the probabilistic representation is singled out as giving the agent's degrees of belief. Additionally, both the structural and substantive axioms on preference seem to conflict with various intuitively permissible sets of preferences, as in

the paradoxes of Allais (1953) and Ellsberg (1961). There is also a vast literature on the apparent inadequacies of the sort of decision theory the representation theorem assumes, independent of a critique of particular axioms. Schick (2003) follows the psychological literature stemming from Kahneman and Tversky (1979) in pointing out the importance of context in decision making (which expected utility theory leaves out), while Nover and Hájek (2004) and Arntzenius et al. (2004) (among others) point out serious problems expected utility theory faces with infinity. Although the theorem doesn't directly assume the existence of fair prices for bets, standard decision theory guarantees that these fair prices exist, so worries about Dutch book arguments can often be translated into worries about decision theory. And finally, representation theorems don't directly put any constraints on the diachronic change of degrees of belief, so they are at best a partial vindication of Bayesianism. However, the fact that decision theory has been such a fruitful field of study in its own right gives representation theorems some additional interest for psychologists, economists, political theorists, and others, independent of their connection to Bayesianism.

1.3. OTHER ARGUMENTS

There are a variety of other arguments that have been given in favor of Bayesianism, to make up for the flaws that these are seen to have. Some are refinements of parts of the other ones, but don't really fit under either category, such as those in van Fraassen (1999) and Greaves and Wallace (2006). But one particular feature that all these arguments seem to share is the fact that they derive rational constraints on degree of belief from apparently pragmatic constraints on preference. This is seen by many to be a failing, since epistemic norms and practical norms seem to be two different sorts of things. Thus, Joyce (1998) seeks 'a non-pragmatic vindication of probabilism' and gives an argument based on accuracy of probabilistic predictions. (Another such argument is discussed in Joyce 2005: 164–5.)

Another non-pragmatic argument that has been quite popular among Bayesian physicists is that given in Cox (1946), and repeated in Chapter 2 of Jaynes (2003). The basic assumptions of this argument are that any set of evidence gives a unique plausibility to any proposition, and that the plausibilities of logically complex propositions are related to those of their constituents by some sort of function. From just these assumptions, one can then prove that the degrees of plausibility must have the same structure as that entailed by the Kolmogorov axioms. Additionally, one can derive some version of the Principle of Indifference (to be discussed more in the section on priors in the next paper), according to which propositions that have equal support from the set of evidence must have equal probability.

However, Halpern (1999) argues that the assumptions of this theorem are far too strong to be supported. The assumption that there is a unique rational set of degrees of belief in any evidential situation is one that even many Bayesians find hard to swallow, but Halpern additionally worries about the infinitary structure required on the space of possibilities for the theorem to go through.

Despite the apparent flaws in every argument for Bayesianism, there are still many Bayesians. One reason is summed up in a quote by Christensen:

Neither [the Dutch book nor the Representation Theorem] comes close to being a knock-down argument for Probabilism, and non-probabilists will find contestable assumptions in both. But each one, I think, provides Probabilism with interesting and non-question-begging intuitive support. And that may be the best one can hope for, in thinking about our most basic principles of rationality. (Christensen 2001: 375)

But another reason many support Bayesianism is because of the many apparent successes of the theory, and its great fruitfulness in dealing with old questions in epistemology and the philosophy of science, as well as in raising new ones. I will discuss these successes, and further questions that they raise, in the next paper.

Short Biography

Kenny Easwaran received his PhD from the Group in Logic and the Methodology of Science at UC Berkeley in May 2008, with a dissertation on the appropriate mathematical formalism for Bayesian conditional probability. He is currently an Assistant Professor at the University of Southern California. He also works on the philosophy of mathematics in addition to formal epistemology.

Notes

* Correspondence: USC School of Philosophy, University of Southern California, Mudd Hall of Philosophy, 3709 Trousdale Pkwy, Los Angeles, CA 90089, USA. Email: easwaran@gmail.com.

¹ He requires that \mathcal{F} be a ‘field of sets’, meaning that it contains E , and is closed under intersection and complement. That is, if A and B are sets in \mathcal{F} , then the intersection $A \cap B$ consisting of all elements that are in both A and B is in \mathcal{F} , and so is the set $E \setminus A$ (also sometimes called \bar{A}) consisting of all elements of E that are not in A . One way to satisfy these conditions is if \mathcal{F} is the power set of E , meaning that it contains every single subset of E . However, when E is infinite, there are fields that leave out some subsets of \mathcal{E} . For instance, the finite subsets together with their complements form a field, as does the collection of all subsets of E definable in any language containing some basic set-theoretic vocabulary. While one may be used to thinking of propositions as arbitrary sets of possibilities, it’s clear that agents often can’t grasp all propositions, and therefore can’t have degrees of belief in them.

² Kolmogorov later considers infinitary generalizations of these axioms, requiring that \mathcal{F} be closed under countably infinite intersections, as well as finite ones, and requiring that P be additive for countably infinite disjoint unions as well as finite ones. These added complications are important for some debates, but inessential for now.

³ When $P_{new}(A) = 1$, this reduces to standard conditionalization, so in that sense, this is a generalization of the standard approach. This approach can also be further generalized to allow a partition of the possibilities into A_1, \dots, A_n rather than just A and $\neg A$. Thus, if A_1, \dots, A_n are incompatible, but their union includes the whole space, then $P(B) = \sum_{i=1}^n P_{old}(B | A_i) P_{new}(A_i)$.

⁴ For instance, if the agent has degree of belief p in A , and q in the disjoint event B , but has degree of belief $p + q - \epsilon$ in the union $A \cup B$, then she will buy a bet on A for $\$p - \epsilon/4$, and buy one on B for $\$q - \epsilon/4$, and sell one on $A \cup B$ for $\$(p + q) - 3\epsilon/4$. The total of the payments gives her a loss of $\epsilon/4$, and the payoffs of the bets will cancel out no matter what happens.

⁵ One might prefer a notion of guarantee that relies on metaphysical or logical necessity, but this makes the problem of logical omniscience discussed in the next paper worse, and will only justify a modified version of the Kolmogorov axioms.

⁶ Footnote 2 of Kyburg (1978) mentions another relevant point, by Teddy Seidenfeld, pointing out that any bookie that has degrees of belief of her own will always prefer one bet in the book to the others, and thus would rather just buy up many of that bet, rather than the whole book.

⁷ Responses to these arguments have been given in Skyrms (1987) and Christensen (1996), among other places, though not everyone finds their responses compelling. The connection between epistemic and practical rationality seems quite different from that discussed in the literature on epistemic contextualism, but perhaps there is some relation.

⁸ Technically, nothing here requires the set of possible outcomes for A to be finite, so these sums should be replaced by appropriate integrals – the details of various versions of these arguments are given in the books cited earlier.

Works Cited

- Allais, M. ‘Le Comportement de l’homme Rationnel Devant le Risque: Critique Des Postulats et Axioms de l’école Americaine.’ *Econometrica* 21 (1953): 503–46.
- Armendt, B. ‘Is There a Dutch Book Argument for Probability Kinematics?’ *Philosophy of Science* 47.4 (1980): 583–8.

- Arntzenius, F. 'Some Problems for Conditionalization and Reflection.' *The Journal of Philosophy* 100.7 (2003): 356–71.
- , A. Elga, and J. Hawthorne. 'Bayesianism, Infinite Decisions, and Binding.' *Mind* 113.450 (2004): 251–83.
- Bradley, D. and H. Leitgeb. 'When Betting Odds and Credences Come Apart: More Worries for Dutch Book Arguments.' *Analysis* 66.290 (2006): 119–27.
- Christensen, D. 'Clever Bookies and Coherent Beliefs.' *The Philosophical Review* 100.2 (1991): 229–47.
- , 'Dutch Book Arguments De-pragmatized: Epistemic Consistency for Partial Believers.' *The Journal of Philosophy* 93.9 (1996): 450–79.
- , 'Preference-Based Arguments for Probabilism.' *Philosophy of Science* 68.3 (2001): 356–76.
- Cox, R. T. 'Probability, Frequency and Reasonable Expectation.' *American Journal of Physics* 14.1 (1946): 1–13.
- Diaconis, P. and S. Zabell. 'Updating Subjective Probability.' *Journal of the American Statistical Association* 77.380 (1982): 822–30.
- Earman, J. *Bayes or Bust?* Cambridge, MA: The MIT Press, 1992.
- Elga, A. 'Defeating Dr. Evil with Self-locating Belief.' *Philosophy and Phenomenological Research* 69.2 (2004): 383–96.
- , 'Reflection and Disagreement.' *Noûs* 41.3 (2007): 478–502.
- , 'Self-locating Belief and the Sleeping Beauty Problem.' *Analysis* 60.2 (2000): 143–7.
- Ellsberg, D. 'Risk, Ambiguity, and the Savage Axioms.' Technical report, The RAND Corporation, 1961.
- Field, H. 'A Note on Jeffrey Conditionalization.' *Philosophy of Science* 45.3 (1978): 361–7.
- de Finetti, B. 'La Prévision: Ses Lois Logiques, Ses Sources Subjectives.' *Annales de l'Institut Henri Poincaré* 7 (1937): 1–68.
- van Fraassen, B. 'Belief and the Problem of Ulysses and the Sirens.' *Philosophical Studies* 77 (1995): 7–37.
- , 'Belief and the Will.' *The Journal of Philosophy* 81.5 (1984): 235–56.
- , 'Conditionalization, a New Argument for.' *Topoi* 18 (1999): 93–96.
- Greaves, H. and D. Wallace. 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility.' *Mind* 115.459 (2006): 607–32.
- Hailperin, T. 'Probability Logic.' *Notre Dame Journal of Formal Logic* 25.3 (1984): 198–212.
- Hájek, A. 'Arguments for—or Against—Probabilism.' *The British Journal for the Philosophy of Science* 59.4 (2008): 793–819.
- Hájek, A. 'What Conditional Probability Could not be.' *Synthese* 137 (2003): 273–323.
- Halpern, J. 'Cox's Theorem Revisited.' *Journal of Artificial Intelligence Research* 11 (1999): 429–35.
- Hitchcock, C. 'Beauty and the Bets.' *Synthese* 139.3 (2004): 405–20.
- Jaynes, E. T. *Probability Theory: The Logic of Science*. Cambridge, UK: Cambridge UP, 2003.
- Jeffrey, R. *The Logic of Decision*. Chicago: University of Chicago Press, 1965.
- Joyce, J. *The Foundations of Causal Decision Theory*. Cambridge, UK: Cambridge UP, 1999.
- , 'How Probabilities Reflect Evidence.' *Philosophical Perspectives* 19 (2005): 153–78.
- , 'A Nonpragmatic Vindication of Probabilism.' *Philosophy of Science* 65.4 (1998): 575–603.
- Kahneman, D. and A. Tversky. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica* 47.2 (1979): 263–91.
- Kolmogorov, A. N. *Foundations of the Theory of Probability*. New York: Chelsea, 1950.
- Kyburg, H. 'Subjective Probability: Criticisms, Reflections, and Problems.' *Journal of Philosophical Logic* 7 (1978): 157–80.
- Lewis, D. 'Sleeping Beauty: Reply to Elga.' *Analysis* 61.271 (2001): 171–6.
- McGee, V. 'An Airtight Dutch Book.' *Analysis* 59.4 (1999): 257–65.
- Meacham, C. 'Sleeping Beauty and the Dynamics of de se Beliefs.' *Philosophical Studies* 138.2 (2008): 245–69.
- Nover, H. and A. Hájek. 'Vexing Expectations.' *Mind* 113 (2004): 305–17.
- Popper, K. *The Logic of Scientific Discovery*, chapter iv*, 326–48. New York: Harper & Row, 1959.
- Ramsey, F. P. 'Truth and Probability.' *The Foundations of Mathematics and other Logical Essays (1931)*. Ed. R. B. Braithwaite. New York: Harcourt, Brace and Company, 1926. 156–98.
- Savage, L. J. *The Foundations of Statistics*. New York: John Wiley and Sons, 1954.
- Schick, F. *Ambiguity and Logic*. Cambridge, UK: Cambridge University Press, 2003.
- Skyrms, B. 'Coherence.' *Scientific Inquiry in Philosophical Perspective*. Ed. Nicholas Rescher. Pittsburgh, PA: University of Pittsburgh Press, 1987. 225–42.
- Teller, P. 'Conditionalization and Observation.' *Synthese* 26 (1973): 218–58.
- Titelbaum, M. 'The Relevance of Self-locating Beliefs.' *The Philosophical Review* 117.4 (2008): 555–606.
- Wagner, C. 'Probability Kinematics and Commutativity.' *Philosophy of Science* 69 (2002): 266–78.
- Weatherston, B. 'Should we Respond to Evil with Indifference?' *Philosophy and Phenomenological Research* 70.3 (2005): 613–35.
- Williamson, J. 'Countable Additivity and Subjective Probability.' *The British Journal for the Philosophy of Science* 50 (1999): 401–16.
- Zynda, L. 'Representation Theorems and Realism About Degrees of Belief.' *Philosophy of Science* 67 (2000): 45–69.