This is a section of doi:10.7551/mitpress/11252.001.0001

The Handbook of Rationality

Edited by: Markus Knauff, Wolfgang Spohn

Citation:

The Handbook of Rationality
Edited by: Markus Knauff, Wolfgang Spohn
DOI: 10.7551/mitpress/11252.001.0001
ISBN (electronic): 9780262366175
Publisher: The MIT Press
Published: 2021

Funding for the open access edition was provided by the MIT Libraries Open Monograph Fund.

The MIT Press

4.1 Subjective Probability and Its Dynamics

Alan Hájek and Julia Staffel

Summary

This chapter is a philosophical survey of some leading approaches in formal epistemology in the so-called *Bayesian* tradition. According to them, a rational agent's degrees of belief—*credences*—at a time are representable with probability functions. We also canvas various further putative "synchronic" rationality norms on credences. We then consider "diachronic" norms that are thought to constrain how credences should respond to evidence. We discuss some of the main lines of recent debate and conclude with some prospects for future research.

1. Outline of the Chapter

You are certain that something exists, slightly less confident that your lottery ticket will lose, less confident again that it will be cloudy tomorrow, still less confident that your lottery ticket will win, and you have no confidence whatsoever that nothing exists. Traditional epistemology traffics mainly in the binary doxastic states of *belief* and *knowledge*, and classical logic is thought by many to provide norms for them (see chapter 3.1 by Steinberger and chapter 5.1 by van Ditmarsch, both in this handbook). But we also have a whole spectrum of belief-like attitudes. Following Ramsey (1926/1990), we may call them *degrees of belief* and regard them as being answerable to a "logic of partial belief": probability theory.[1] This, in turn, gives that mathematical theory one of its main interpretations: *subjective probability*.

The theory of subjective probability provides us with a *normative*, *numerical* model of degrees of belief. It uses precise numbers between 0 and 1 to represent the degree to which someone believes something. Moreover, it formulates normative requirements that their degrees of beliefs have to obey in order to be perfectly rational. Hence, the philosophical investigation of subjective probability is less concerned with the question of whether people are in fact rational—this question is delegated to empirical psychologists (see chapter 4.5 by Chater & Oaksford, this handbook)—and more concerned with characterizing norms of ideal rationality and reasoning. In this chapter, we survey a variety of purported rational requirements and the philosophical justifications that have been offered in their support—both *synchronically* (for an agent at a given time) and *diachronically* (for an agent at different times).

2. The Probability Calculus

To understand subjective probability, we must first understand *probability*. Kolmogorov's axiomatization (1933/1950) remains the orthodox formalization of probability—the so-called *probability calculus*.[2] We begin with a "universal set" Ω regarded as the set of all possibilities of interest or of all possible outcomes of a given random experiment. Probabilities are numerical values assigned by a real-valued function P to certain subsets of Ω—*events*—obeying the following axioms:

1. For all events X, $P(X) \geq 0$. (Nonnegativity)
2. $P(\Omega) = 1$. (Normalization)
3. For all disjoint events X and Y, $P(X \cup Y) = P(X) + P(Y)$. (Finite Additivity)

These axioms are often reformulated with events replaced by sentences in a formal language. Normalization becomes

$P(T) = 1$, where T is any tautology.

And Finite Additivity becomes

for all logically incompatible sentences X and Y, $P(X \vee Y) = P(X) + P(Y)$.[3]

Conditional probability is probability *relative to*, or *given*, some body of evidence or information. Following Kolmogorov, we symbolize "the probability of X, given Y" as $P(X \mid Y)$ and define it as

$$P(X \mid Y) = \frac{P(X \cap Y)}{P(Y)}, \quad \text{if } P(Y) > 0.$$

For more on the probability calculus, see Hájek and Hitchcock (2016).[4]

This formalism is brought to life when we interpret "*P*." It might be regarded, for example, as representing *objective chances* in the world, independent of anyone's opinions; these might be understood as relative frequencies, as propensities, or as falling out of the best systematization of the universe—see Hájek (2019) for more discussion and references. However, according to the interpretation of interest here, "*P*" represents a rational agent's degrees of belief, also known as *degrees of confidence*, or *credences*. For example, your credence that a roll of a fair die lands an even number should be

$$P(2 \cup 4 \cup 6) = P(2) + P(4) + P(6) = \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = \frac{1}{2}.$$

And your credence that it lands 6, *given* that it lands an even number, should be

$$\frac{1/6}{1/2} = \frac{1}{3}.$$

Going in the other direction, we might begin by assuming that a rational agent has various degrees of belief and ask how to represent them formally. The standard view among philosophers, economists, and others in the "Bayesian" tradition is that someone's degrees of belief should be representable by a probability function conforming to this calculus. This view is known as *probabilism*. This requires some analysis: *what are* degrees of belief? And it requires some justification: *why should* they obey these axioms? The justification will be tailored to the analysis.

3. Accounts of Subjective Probability, and Arguments for Probabilism

3.1 The Betting Account

The dominant approaches see degrees of belief as intimately tied to action and decision making. A behaviorist approach identifies degrees of confidence with betting behavior (e.g., de Finetti, 1937/1964). Its gist is as follows:

> Your probability of E is x iff x is the price for which you would either buy or sell a bet that pays a unit of money if E and that pays nothing otherwise.

There are various problems with this account. It assumes that there is a unique such price when there may be none or more than one. Moreover, the account is implausible when E is an event under your control. For example, you assign low probability to your impersonating a chicken today, but that may change if your doing so wins a $1 bet. (If $1 does not suffice, make it $1,000,000—raising the further problem that your betting behavior is sensitive to the stakes involved, while your credences are not.) More generally, measuring your credences with bets may change what those credences are, as Ramsey observed. (Indeed, someone offering a high-stakes bet on a proposition may make you suspicious that they have inside knowledge about it!) And a given betting price may be jointly determined by many components—the value you attach to money, the thrill of gambling, a public show of bravado, and so on—only one of which is your credence.

3.2 The Dutch Book Argument

Despite its limitations, many authors regard the betting account as "fundamentally sound," in Ramsey's (1926/1990) words. As such, it is appealed to in one of the most important arguments for probabilism: the *Dutch Book argument*. A *Dutch book* is a set of bets bought or sold at such prices as to guarantee a net loss. We have the celebrated *Dutch book theorem*, first stated by Ramsey, and proven by de Finetti (1937/1964):

> If a set of (normalized) betting prices violates the probability calculus, then there is a Dutch book consisting of bets at those prices.

So, if your betting prices violate the probability calculus, you are susceptible to a Dutch book. However, the argument continues, if you are susceptible, then you are irrational. Hence, rationality requires your betting prices to obey the probability calculus. Identifying your credences with your betting prices, the argument concludes: rationality requires your credences to obey the probability calculus.

The argument is often personified with a canny bookie fleecing a hapless agent who has nonprobabilistic credences. For example, suppose you assign probability ½ to the Democrats winning the next election, probability ½ to the Republicans winning, but only probability 0.9 to either the Democrats or the Republicans winning. The bookie could sell you for 50 cents a bet that pays $1 if the Democrats win, sell you for 50 cents a bet that pays $1 if the Republicans win, and buy from you for 90 cents a bet that pays $1 if either of them win. Initially, you pay him a total of $1 and he pays you 90 cents, with you losing 10 cents upfront. Then, come what may, the bets that you hold pay exactly the same as the bet that he holds, so that your net loss remains the same.

Often neglected, but equally important, is the *converse* Dutch book theorem (Kemeny, 1955):

> If a set of betting prices (normalized to 1) *obeys* the probability calculus, then there is *no* Dutch book consisting of bets at those prices.

This dispels the concern that even so-called *coherent* agents, whose credences obey the probability calculus, might be susceptible to Dutch books. Not so: coherence inoculates us!

The Dutch book argument is often criticized for bringing out a *pragmatic* rather than an *epistemic* defect in an incoherent set of credences. Furthermore, some critics complain that an incoherent agent need not fall into the trap of being Dutch-booked—for example, you might just walk away instead of making the sure-loss package of bets. In response, Skyrms (1980), following Ramsey, interprets the argument as merely dramatizing an inconsistency in such an agent's evaluations—an epistemic defect. For more discussion, see Skyrms (1993), Vineberg (2001, 2016), and Hájek (2009).

3.3 Representation Theorems

We have noted some problems with the betting account of credences, which carry over to the Dutch book argument: if your credences come apart from your betting prices, a defect in the latter does not entail a defect in the former. Let's move on, then, to another account of credence and another kind of argument for probabilism that is based on it. We have Ramsey to thank for both.

He presupposes an operationalist attitude to defining such a theoretical term: "the degree of a belief . . . has no precise meaning unless we specify more exactly how it is to be measured" (Ramsey, 1926/1990, p. 167). His guiding idea is that probability "is a measurement of belief . . . *qua* basis of action" (Ramsey, 1926/1990, p. 171) and to use an agent's preferences among gambles to provide this measurement. He proposes a number of axioms governing such preferences. Some are "consistency" assumptions, such as transitivity, while others are structural, aimed at streamlining the mathematics. Together, they allow him to calibrate a *utility* scale for the agent, a measure of how desirable various gambles and their outcomes are. This, in turn, determines her credences as ratios of differences of utilities. He goes on to add further axioms that guarantee their obeying the probability calculus.

Savage (1954) similarly derives probabilities and utilities from preferences that are constrained by certain axioms. However, he is more concerned with the *normative* status of the axioms than Ramsey is: they constrain *rational* preferences (see also chapter 8.1 by Grüne-Yanoff, this handbook). Savage proves a *representation theorem*: for a given (representation of a) set of such preferences, there is a class of utility functions, unique up to positive linear transformations, and a unique probability function that jointly "represent" the preferences in the following sense: the *expected utility* of an option ("act") is the weighted average of the utilities of its possible outcomes associated with a set of "states," the weights provided by the probabilities of the states. A utility function (chosen from such a class) and probability function are said to jointly *represent* an agent's preferences when she prefers *A* to *B* iff the expected utility of *A* is greater than that of *B*. Jeffrey (1965) offers an alternative axiomatization of rational preferences to Savage's, and Bolker proves another representation theorem (cited by Jeffrey), which typically does not yield a unique probability function.

Such representation theorems undergird an argument for probabilism: any rational agent has preferences that obey the axioms; as such, she is representable as an expected-utility maximizer, with weights interpreted as her credences. These credences obey Kolmogorov's axioms. However, some question the normative status of some of the preference axioms—they are there more to facilitate the mathematical representation. Others question the step from the rational requirement of the agent being *representable* probabilistically to the rational requirement of her credences *being* probabilities. For example, she may *also* be representable *non*probabilistically (see Zynda, 2000). Indeed, she may be representable in entirely implausible ways—for example, as a voodoo spirit maximizer (see Hájek, 2008).

Representation-theorem approaches to credences allow and even support an anti-realist stance regarding them: they are merely artifacts of a representation. (Stefánsson [2017] defends this view.) By contrast, a more realist understanding of credences fits well with our next arguments for probabilism. (See Eriksson & Hájek [2007] for further discussion of the metaphysics of credences, including their view that they should be regarded as primitive, irreducible to anything else; see also note 2 for references to approaches that take comparative probabilities as primitive.)

3.4 Accuracy Arguments

Belief is typically regarded as being a genuine mental state rather than merely a representational artifact. Moreover, it is often said to have a constitutive aim: *truth*. As such, a given belief may be evaluated on whether or not it achieves this aim—whether it is true or false (see also chapter 3.1 by Steinberger, this handbook). However, intermediate credences apparently *cannot* be said to be true or false. Are they nonetheless answerable to a truth norm?

A thinker may assign intermediate degrees of confidence to strike a balance between the competing aims of believing truths and avoiding belief in falsehoods. Ideally, the agent assigns full credence to all true propositions and zero credence to all false propositions. When she assigns intermediate credences, we can measure how close they are to the truth. A specific class of measures, called *scoring rules* or *accuracy measures*, has been developed for this purpose.

A popular measure is the Brier score: if the agent's credence in some event E is x, and E actually occurs, then the Brier score for this credence is $(1 - x)^2$; if E doesn't actually occur, the score is $(0 - x)^2$. We can compute the Brier score for an entire credence function at a particular world by summing the individual scores for the agent's credences at that world. The closer the score is to zero, the better, since a lower score indicates better accuracy or greater closeness to the truth. The Brier score has been put to practical use, for example, to score the accuracy of weather forecasters and political analysts. A feature that makes the Brier score, among other scoring rules, particularly suited for this purpose is that it is *strictly proper*: someone who is asked to report her credences expects to minimize the score she will receive by reporting her actual credences. By contrast, if we used an improper scoring rule, such as the absolute distance measure $|1 - x|$ or $|0 - x|$, to score the agent's credences' distance from the truth, she would sometimes be better off lying about her credences to improve her (own expectation of her) accuracy score.

Suppose we adopt the Brier score to measure accuracy. Then, if an agent has nonprobabilistic credences, there is an alternative probabilistic credence function that is more accurate in every possible world. Conversely, if she has coherent credences, there is no alternative probabilistic credence assignment that is more accurate in every possible world. Hence, all and only coherent credence functions avoid being *accuracy dominated*. A version of this mathematical result was first demonstrated by de Finetti (1974, pp. 87–91) and has since been developed further to show that it holds for additional suitable scoring rules besides the Brier score. The theorem is the central premise of the *accuracy dominance argument* for probabilism, which claims that having accuracy-dominated credences is a rational defect, which can only be avoided by having probabilistic credences (for developments of the theorem and the argument, see Joyce, 1998; Leitgeb & Pettigrew, 2010a, 2010b; Pettigrew, 2016).

The accuracy dominance argument is popular because it doesn't rely on pragmatic considerations—it is thought to be an argument for epistemic rationality that proceeds from purely epistemic premises. However, for the argument to be sound, we need to select a specific accuracy measure from the class of eligible strictly proper scoring rules (Bronfman, manuscript). This has presented something of a problem, since it is difficult to find compelling epistemic grounds for selecting a specific accuracy measure.[5]

4. Other Putative Synchronic Norms

4.1 Regularity

"If it can happen, then it has some probability of happening." Many authors have offered some version of this intuitive slogan as a constraint on credences (e.g., Carnap, 1950; Lewis, 1980; Shimony, 1955). More formally, *regularity* constraints have the form

If X is possible, then $C(X) > 0$,

where C represents the credences of some rational agent. There are various candidates for the sense of "possibility" here, but *logical*, *metaphysical*, and *epistemic* possibility are especially prominent.

There are both theoretical and pragmatic arguments in favor of regularity. Theoretical: Your credences should reflect your evidence; if your evidence does not decisively rule out X, so that X remains epistemically possible, your credence should not rule it out either. Furthermore, such an open mind regarding X is required for you to be able to learn it (Lewis, 1980). Pragmatic: If you assign probability 0 to a possible proposition X, then you are susceptible to a *semi-Dutch Book* (Shimony, 1955): you will consider fair a bet that you could lose and that could not possibly win you any money. According to the betting interpretation, you are prepared to pay \$1 for a \$1 bet on X. This seems irrational, since at best you will break even, and you could lose \$1. However, see Hájek (2012) for rebuttals of these arguments (and others).

Furthermore, there are both theoretical and pragmatic arguments *against* regularity. Theoretical: If the space of possibilities is sufficiently large, regularity is mathematically impossible. For example, if there are uncountably many possible outcomes, then a (real-valued) probability function cannot assign positive probabilities to all of them. Indeed, uncountably many of its assignments must be 0 (see Hájek, 2003). Pragmatic: Assuming that your gaining infinite expected utility is possible (as in Pascal's wager and the St. Petersburg paradox[6]), regularity requires you to assign positive probability to your gaining it. This, in turn, requires you to value this prospect infinitely, however small this probability—the same as you would if you had it for sure (see Hájek, 2012; for further opposition to regularity, see also Easwaran, 2014; Levi, 1989; T. Williamson, 2007).

4.2 The Principal Principle

We have mentioned that the "*P*" that Kolmogorov axiomatized might be interpreted objectively as *chance*. Lewis offers a further rationality constraint on credences, which codifies a certain kind of alignment between them and chances. Here is his classic statement of it (Lewis, 1980, p. 266):

> **The Principal Principle.** Let *C* be any reasonable initial credence function. Let *t* be any time. Let *x* be any real number in the unit interval. Let *X* be the proposition that the chance, at time *t*, of *A* equals *x*. Let *E* be any proposition compatible with *X* that is admissible at time *t*. Then
>
> $C(A/XE) = x.$

An initial credence function is a "prior" probability function, representing one's credences before one receives *any* evidence. An admissible proposition is one "whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes" (Lewis, 1980, p. 272). Typically, information about the past is admissible, whereas information about how the chance process turns out is inadmissible. An instance of the Principal Principle for a toss of a coin after noon is:

$C(\text{heads} \mid chance_{\text{noon}}(\text{heads}) = ½$, and another toss of the coin landed heads at 11:59 am$) = ½$.

Lewis's talk of an "initial" credence function may come as a surprise—it is not *your* credence function or that of any rational agent who has learned *anything*. Lewis (1994) presents the Principal Principle differently, speaking instead of "a rational credence function for someone whose evidence is limited to the past and present—that is, for anyone who doesn't have access to some very remarkable channels of information" (p. 483). This far more natural (nonequivalent) formulation is how the principle is usually understood. It is also surprising that the vast literature on the Principal Principle overlooks the change in formulation.[7] Pettigrew (2016) gives an accuracy-based justification for the principle, as well as a comprehensive overview and discussion of different formulations of it. For a Dutch book–style argument, see Pettigrew (2018).

4.3 The Reflection Principle

The Principal Principle is sometimes called an *expert principle* or a *deference principle*, because it requires rational thinkers to treat a source of probabilities—the chance function—like an expert and to defer to its assignments. Another such principle is the *Reflection Principle* (van Fraassen, 1984), which enjoins thinkers to defer to their own future credences. Where *t* is a particular time, and $t + n$ is a later time,

$$C_t(A \mid C_{t+n}(A) = x) = x.$$

Informally stated: Given that at a later time your credence in some claim *A* will be *x*, your credence in *A* should already be *x* at the current time. For example, suppose you are certain that once you start your car tomorrow morning and hear the rattling of its old engine, you will be highly confident that the car needs servicing; then you should be just as confident *now* that the car will need servicing tomorrow.

Van Fraassen argued for the principle by showing that there is a Dutch book argument for it similar to the one for conditionalization (more on that below; for discussion, see also Briggs, 2009; Christensen, 1991; Mahtani, 2015). There is also an argument showing that obeying Reflection maximizes the expected accuracy of one's credences (Easwaran, 2013). However, it has been pointed out that, for the principle to be plausible, the range of situations in which it applies must be restricted. If you expect your future self to be either irrational or forgetful, then you should not defer to your future credences (Arntzenius, 2003; Christensen, 1991; Talbott, 1991). Van Fraassen (1989, 1999) offers a modified formulation of the principle to address these types of worries.

4.4 The Indifference Principle

Another putative norm of rationality is the *Indifference Principle*, which applies when we either have no evidence, or our evidence does not discriminate between the relevant possibilities. Here's a version of it:

> If there are *n* possibilities, and an agent has no evidence that discriminates between them, then the rational credence for her to assign to each possibility is $1/n$.

For example, suppose you are informed that you have been entered into some kind of lottery alongside 99 other participants. You know nothing about how the winner is determined. The Indifference Principle then advises that you should assign a 1/100 credence to winning.

The indifferent credence distribution is supported by minimax reasoning about the accuracy of one's credences: it minimizes the worst-case inaccuracy of one's credences (Pettigrew, 2016). It also minimizes *information* among all distributions over the possibilities. For a discrete probability distribution $P = (p_1, p_2, \ldots, p_n)$, the *information* of *P* is defined as

$$\sum_i p_i \log_2 p_i$$

(where $0 \log_2 0$ is regarded as 0). Equivalently, the indifferent distribution maximizes *entropy*, defined (by Shannon, 1948) as

$$-\sum_i p_i \log_2 p_i.$$

Entropy is regarded as a measure of the uncertainty of a distribution. A distribution that concentrates probability 1 on a single possibility is maximally informative and minimally uncertain; a uniform distribution is minimally informative and maximally uncertain. Entropy may also be generalized to continuous probability distributions, and the Indifference Principle may be generalized to the principle that the rational probability assignment, subject to a given constraint (e.g., background knowledge), is the distribution with maximum entropy that meets the constraint. Jaynes (2003) is a prominent advocate of this principle.

The Indifference Principle—and its generalization—is controversial, however, because it sometimes gives conflicting recommendations, made vivid in Bertrand's paradoxes (Bertrand, 1889). Van Fraassen (1989) illustrates this problem. Consider a factory that produces cubes with side lengths up to 1 foot. What should your credence be that a randomly selected cube has a side length up to ½ foot? The natural answer is of course ½. But this assumes that side length is the relevant quantity over which we should uniformly distribute our credence. We might just as well take the cubes' *volume* to be this quantity. The factory produces cubes with volumes up to 1 cubic foot. Then the question becomes how confident we should be that a randomly selected cube has a volume up to $½^3 = ⅛$ cubic foot. Now the natural answer is ⅛! And we get yet another answer if we consider the area of the cube's faces instead. Since none of these ways of dividing up the possibilities into symmetrical alternatives stands out as privileged, the Indifference Principle cannot deliver a unique answer. It's a matter of debate whether this is a fatal objection to the principle. See Norton (2008), White (2009), and Novack (2010) for some defenses of the principle.

5. Diachronic Norms

5.1 Conditioning

Conditionalization or *conditioning* is the most widely endorsed rule that governs how a thinker should respond when she becomes certain of a new piece of evidence:

> When a thinker learns a piece of evidence E with certainty (and learns nothing in addition), then her new credence in any proposition A is equal to her old conditional credence $C(A|E)$.

For example, suppose you are 90% confident that your neighbor is home, given that his lights are on. Otherwise, you are only 50% confident that he is home. If you then see that his lights are on, your confidence that he is home should be 90%.

Both Dutch book and accuracy arguments have been offered in support of conditioning. The Dutch book is *diachronic*, because it requires that bets be placed at two different times—before and after learning the new evidence. Lewis (1999) shows that violations of conditionalization lead to Dutch book vulnerability. Skyrms (1987) provides the needed converse theorem: conditionalizers are immune to Dutch books (provided they don't commit any other Dutch-bookable offenses). Critics worry that diachronic Dutch book arguments are less compelling than synchronic ones. Christensen (1991) argues that Dutch-bookable inconsistencies between one's earlier and later self's credences are no more irrational than similarly Dutch-bookable disparities between one's own credences and one's spouse's credences. Taking the betting story literally, Levi (1987) argues that diachronic Dutch books are ineffective, since the agent would simply refuse the initial bets if she were aware of the bookie's strategy. Skyrms (1993) shows that once the decision faced by the agent is fully specified, she can be enticed to bet—and lose—even if she knows what the bookie is up to.

There are also variations of expected accuracy arguments for conditioning: relative to her initial credences, she will maximize her credences' expected accuracy upon learning whether E iff she conditionalizes (Easwaran, 2013; Greaves & Wallace, 2006; Leitgeb & Pettigrew, 2010b; Pettigrew, 2016; Schoenfield, 2017).

A concern about the accuracy arguments for conditioning is that they don't truly show that an agent who does otherwise is irrational. The credences C that recommend conditioning are those the agent accepted before she knew whether E, and the conditioning strategy recommends C only before it is settled whether E. Once she has learned whether E, she recognizes that C is outdated. Why should she update in a way that's recommended by credences that are then known to be defective (Pettigrew, 2016)? Moreover, as long as she adopts any coherent credences after learning whether E, any strictly proper inaccuracy measure will underwrite the result that the agent expects her current credences to be more accurate than any alternative credences (Easwaran, 2013). Pettigrew (2016) thus contends that these types of arguments can only show that rational agents must *plan* to conditionalize but not that they must follow through.

A different justification for conditioning comes from "minimal mutilation" considerations. Start with a credence function C, and impose the constraint that E is to be assigned probability 1 (and that it is the strongest such proposition). We want to move to the probability function *closest* to C that meets the constraint. There are various "distance" measures on the space of probability functions. A notable one is the *Kullback–Leibler divergence*, or *relative entropy*. For discrete distributions $P = (p_1, p_2, \ldots)$ and $Q = (q_1, q_2, \ldots)$, it is defined as

$$D(P, Q) = \sum_i p_i \log_2\left(\frac{p_i}{q_i}\right).$$

It can be shown that $C(\cdot \mid E)$ is the function that minimizes entropy relative to C among all functions that meet the constraint. Other distance measures that underwrite conditioning in this way are the *Hellinger distance* and the *variation distance* (see Diaconis & Zabell, 1982). These results support conditioning insofar as they show that an agent who follows it changes her credences conservatively: she doesn't move away from her original credences any more than the evidence warrants, in a certain sense. However, one may ask why these particular distance measures are epistemically significant, since other distance measures deliver results that can disagree with conditioning. Moreover, one may ask why it's important to stay close to one's previous credences when they have become outdated.

5.2 Jeffrey Conditioning

One limitation of conditioning is that it applies only in cases where a thinker learns a piece of evidence with certainty. But sometimes we gather information without becoming certain of it. For example, suppose you glance at a tablecloth in poor light. Based on what you saw, you become more confident that the fabric is blue but not 100% confident. In this type of case, it is controversial whether there is any proposition that you have learned with certainty (although see Skyrms, 1987).

Jeffrey conditionalization, or *Jeffrey conditioning*, is an updating rule (named after Jeffrey, 1965) that can incorporate uncertain learning. Suppose $\{E_1, \ldots, E_n\}$ is a partition—a set of mutually exclusive and jointly exhaustive propositions—such that the agent's experience will change her credences in the elements of the partition from $C_{\text{old}}(E_i)$ to $C_{\text{new}}(E_i)$. Then, Jeffrey conditioning is this updating rule:

$$\text{for any } A,\ C_{\text{new}}(A) = \sum_i C_{\text{old}}(A \mid E_i) C_{\text{new}}(E_i).$$

In words: one's new credence in A is the weighted sum of one's old conditional credences $C_{\text{old}}(A \mid E_i)$, where the weights are one's new credences in the partition members. When the thinker becomes certain of a particular E_i, Jeffrey conditioning reduces to standard conditioning. One interesting difference between standard and Jeffrey conditioning is that the former is commutative, while the latter is not. This means that when one receives more than one piece of evidence, standard conditioning always delivers the same credences in the end, regardless of the ordering of the updates. By contrast, with Jeffrey conditioning, the resulting credences need not be the same if the updates are made in a different order (Diaconis & Zabell, 1982; Levi, 1967).

There is currently no accuracy-based argument in favor of Jeffrey conditioning, but there is a diachronic Dutch book argument for it (Armendt, 1980; Skyrms, 1987). Additionally, it can be shown that the credence function that is arrived at by Jeffrey conditioning is the credence function *closest* to the agent's original credences that satisfies the constraint of assigning the values $C_{\text{new}}(E_i)$ to the members of the given partition, according to various "distance" measures (Kullback–Leibler, Hellinger, variation distance; see Diaconis & Zabell, 1982).

5.3 Relative Entropy, and More General Dynamics

The general problem for probability dynamics is this. Start with a credence function C, and impose a constraint K; to which credence function meeting K should C be revised? We have just seen solutions to this problem for the particular constraints that E be assigned credence 1 (conditioning) and that specified credences be assigned across a partition (Jeffrey conditioning). But we can imagine various other such constraints—for example, assign specified *conditional* credences across a partition or assign a probability in the interval $[x, y]$ to E. Any "distance" measure gives us a candidate for answering this general problem: move to the function closest to C, by that measure's lights, that satisfies K (when there is one). Relative entropy is an especially popular such measure (although van Fraassen [1981] raises his "Judy Benjamin problem" against it for a conditional-credence constraint).

6. Uniqueness

The putative norms of synchronic and diachronic rationality that we have specified so far don't generally single out a unique rational credence function for an agent at a time. Some philosophers, usually called *permissivists*, think that for any given body of evidence, there is more than one overall doxastic state a rational thinker could adopt on its basis (see, e.g., Douven, 2009; Kelly, 2013; Meacham, 2014; Schoenfield, 2014; for further references, see the excellent survey of the debate in Kopec &

Titelbaum, 2016). Proponents of *Uniqueness* disagree. Proponents of *Intra*personal Uniqueness argue that, for any given thinker who possesses a particular body of evidence, there is only one rational overall doxastic state this thinker can adopt in response to her evidence. Proponents of *inter*personal uniqueness argue for the even stronger thesis that there is only one rational response to any given body of evidence, the same for every thinker. (Early defenses of versions of Uniqueness can be found in Christensen, 2007; Feldman, 2007; White, 2005; and, more recently, in Dogramaci & Horowitz, 2016; Greco & Hedden, 2016.[8]) However, proponents of Uniqueness must admit that even if they are right, we probably won't ever know what these uniquely rational doxastic attitudes are.

7. Relaxing the Idealizations

The Bayesian theory of rational credences is highly idealized both in the manner in which it represents credences and in the normative demands it makes on them. Though we have so far granted that philosophers are mainly concerned with characterizing norms of *ideal* rationality, we may legitimately ask how that project relates to norms of rationality for *non-ideal* agents like us.

7.1 Imprecise Probabilities

One way in which the Bayesian view is idealized is in its representation of credences as being precisely point-valued. But realistically, many of our credences are not nearly so specific, and hence representing them this way is a distorting simplification of reality. Moreover, it has been argued that when our evidence is very unspecific, it would in fact be unfitting to form a precise credence on its basis (see, e.g., Joyce, 2005; Levi, 1985). For example, what is your credence that it's cloudy in Vladivostok? It seems hard to come up with a precise value. In response, it has been suggested that we represent credences as *interval-valued*. For example, one's credence that it's cloudy in Vladivostok might be represented as [0.1, 0.9]. More formally, and more generally, the entirety of someone's potentially imprecise credences could be represented as a set of precise credence functions that together capture the person's imprecise epistemic state.

However, this approach still has the problem that the intervals have artificially precise endpoints: where previously there was one precise value, now there are two! Furthermore, once people's credences are represented as sets of probability functions, the Bayesian norms of synchronic and diachronic rationality, and the arguments supporting them, have to be reformulated accordingly, which is not always straightforward. There is also a substantial debate about how decision theory should work when we allow imprecise credences, with Elga (2010) raising a difficulty (see Rinard, 2015, for a response).

7.2 Forgetting/Information Loss

Another highly idealized assumption made by Bayesian representations of people's credences is that thinkers never forget or lose information. When an agent learns a piece of evidence with certainty, standard Bayesian models require that she remains certain of it. This is of course psychologically highly unrealistic, and it is also not obvious that forgetting or losing information is a rational defect. Titelbaum (2012) offers an account of how standard Bayesian updating rules need to be adapted in order to accommodate information loss (for important examples, see also Arntzenius, 2003) and self-locating credences (concerning who one is, where one is, and what time it is).

7.3 Human Limitations in Approximating Ideal Rationality

Even if we modify the standard Bayesian theory to allow for imprecise credences and information loss, the requirements of ideal rationality are still so demanding that humans cannot be expected to fully comply with them. Most obviously, the requirement that one must assign every tautology credence 1 exceeds the reasoning capacities of humans. Psychologists have also found evidence for other deviations from Bayesian rationality, such as our tendency not to change our credences enough in response to incoming evidence (Slovic & Lichtenstein, 1971; see also chapter 8.3 by Glöckner and chapter 8.5 by Hertwig & Kozyreva, both in this handbook). There is a major debate in psychology about how to interpret the relevant data and about whether such results demonstrate that humans are by and large rational (because they often come close to complying with Bayesian norms) or irrational (because they standardly deviate from them).

Once we grant that full compliance with Bayesian norms is excessively difficult for us to achieve, we may regard the norms as regulative ideals that we should approximate as much as is feasible. This raises two questions: First, what exactly does it mean to approximate ideal Bayesian rationality? Second, why is it beneficial to become less irrational, even if we can never reach the ideal?

A good way to measure how well someone's irrational credences approximate ideal rationality is to treat a credence function as a vector and to measure its distance from the closest vector that represents a rational credence function. For example, suppose someone has the following incoherent credences: $C(p) = 0.3$, $C(\sim p) = 0.4$, in vector notation: $\langle 0.3, 0.4 \rangle$. Different distance measures will identify different credence functions as closest. For example, we may use squared Euclidean distance as our distance measure, which is defined for vectors $X = \langle x_1, \ldots, x_n \rangle$ and $Y = \langle y_1, \ldots, y_n \rangle$ as

$$d(X,Y) = (x_1 - y_1)^2 + (x_2 - y_2)^2 + \cdots + (x_n - y_n)^2.$$

Then the closest coherent credence function is $C^*(p) = 0.45$, $C^*(\sim p) = 0.55$, and the distance between C and C^* is 0.045. But if we use other measures, such as absolute distance or the Kullback–Leibler divergence, we'll get different results (e.g., for the latter, the closest coherent credence function is $C^*(p) = 0.438$, $C^*(\sim p) = 0.572$). Which distance measure we use becomes important when we wonder why it is beneficial for thinkers to become more coherent, even if they can't become perfectly coherent. Suppose a thinker reduces the incoherence in her credence function by replacing her original incoherent credence function C_1 with a less incoherent credence function C_2 that is on the direct path toward the closest coherent credence function, as measured by squared Euclidean distance. C_2 is better than C_1 in at least two respects: C_2 is more accurate than C_1 in every possible world, according to the Brier score. Moreover, C_2 is vulnerable to a lower guaranteed loss from a Dutch book than C_1, provided we standardize the size of the bets the bookie is allowed to use. But what if we had used a different measure of closeness to coherence? For example, for the Kullback–Leibler measure, the accuracy result just mentioned holds only if we adopt the log-inaccuracy measure but not if we use the Brier score (see De Bona & Staffel, 2017, 2018; for earlier approaches to the problem, see also Schervish, Seidenfeld, & Kadane, 2002; Zynda, 1996). Hence, we can explain why it is good for thinkers to become less incoherent in specific ways, even if they can't ever have fully coherent credences: it gets them a greater portion of the accuracy benefits and the Dutch book–avoidance benefits that come with perfect coherence. However, these results crucially depend on finding the right combinations of distance measures, Dutch book, and inaccuracy measures.

7.4 Disagreement and Higher-Order Evidence

The fact that we can learn about our own imperfections also raises the question of how we should respond to information about our own ability to correctly assess evidence. Cases in which we disagree with others can provide us with such higher-order evidence. Suppose you and I have agreed to evenly split our lunch bill. We both do the calculations but get different answers. One (or both) of us must have made a mistake (even by permissivist lights). What's the best repair strategy? *Conciliationists* maintain that we should both lower our credence that our respective answers are correct (Christensen, 2007). *Steadfasters*, by contrast, claim that if I've in fact assessed my evidence correctly, higher-order evidence of this type should not move me (Kelly, 2005).[9] Of course, not every disagreement indicates that someone has made a mistake. Different scenarios require different kinds of updates: if you disagree with me because you have more evidence, I should defer to you, but not if you have less evidence. If we disagree because we started from different permissible prior probabilities—before the evidence came in—perhaps we can both stick to our guns. If we find out we independently arrived at the same (high) degree of confidence in something, this might warrant raising our confidence even more (Lasonen-Aarnio, 2013). Claims about how rational agents should respond to disagreement and higher-order evidence thus yield further constraints on which prior probabilities are rationally permissible.

8. Conclusion

Formal epistemology has been one of the greatest growth areas in philosophy in the past couple of decades, and Bayesian approaches have been among the most dominant in formal epistemology. A very simple axiomatization codifies the putative synchronic norms on credences, and simple updating rules codify the putative diachronic norms on them. The result is an elegant theory that allows rigorous proofs of various claims that are taken to have epistemic significance. No wonder it has become so popular!

Moreover, there is still plenty more work to be done. Table 4.1.1 roughly represents the current state of play in Bayesian epistemology. The cells with checkmarks represent the putative norms that have already been given a particular type of putative justification (as far as we are aware at the time of writing). An empty cell can be empty for one of two reasons: either the missing justification is impossible, or it is just a matter of historical contingency that nobody has yet proved the requisite theorem; indeed, in some cases, perhaps nobody has even tried. The corresponding cells may thus be ripe for the taking. Moreover, when additional norms and justification strategies are proposed, the table itself will grow.

Then there are various other state-of-the-art topics that we have only been able to discuss briefly here: Uniqueness; whether credences may, or even must, be imprecise; the psychological reality (or otherwise) of Bayesianism

Table 4.1.1
Summary of putative justifications of norms

Putative justifications → Putative norms ↓	Dutch book	Accuracy	Representation theorem	Maximal entropy	Minimal distance	. . .
Probabilism	✔	✔	✔			
Regularity	✔					
Principal Principle	✔	✔				
Reflection Principle	✔	✔				
Indifference Principle		✔		✔		
Conditionalization	✔	✔			✔	
Jeffrey Conditionalization	✔				✔	
. . .						

and the extent to which humans are answerable to its norms; disagreement; and so on. And finally, there are still other "hot" topics that we have not even touched on: other putative norms; the relationship between quantitative credences and binary beliefs (which are a stock-in-trade of traditional epistemology); the relationship between credences and knowledge (likewise); quantitative versus comparative credences; Bayesianism versus other representations of uncertainty (Dempster–Shafer functions, ranking functions, belief revision theory, possibility theory, etc.); Bayesian decision theory—or rather, rival formulations of it; and more. We may thus be highly confident that Bayesian epistemology will remain a fertile research program for a long time to come!

Acknowledgments

We thank Arif Ahmed, Luc Bovens, Nicholas DiBella, Edward Elliott, Orri Stefánsson, Jeremy Strasser, Timothy Luke Williamson, James Willoughby, and the editors for very helpful comments.

Notes

1. It's an interesting question how we might go from the kinds of intuitive (nonnumerical) attributions of confidence with which we began to *numerical* degrees of belief. See, for example, Fishburn (1986) for discussion of the relationship between comparative and numerical probabilities. Stefánsson (2017) argues that comparative probabilities are primitive and psychologically real but that numerical probabilities are neither.

2. For alternative axiomatizations that take conditional probability to be primitive, see Popper (1959) and Rényi (1970).

3. Kolmogorov goes on to give an infinitary extension of the Finite Additivity axiom:

3′. (Countable Additivity) If $\{A_i\}$ is a countably infinite collection of (pairwise) disjoint events, then

$$P\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} P(A_n).$$

However, the status of this axiom is more controversial (see, e.g., de Finetti, 1974), and we will not assume it in what follows. Moreover, a sentential reformulation of Countable Additivity requires some finessing, as classical logic admits only finite disjunctions.

4. It is a difficult question how conditional probabilities are related to probabilities of conditionals. We refer interested readers to chapter 4.4 by Pfeifer, chapter 4.6 by Oberauer and Pessach, and chapter 6.2 by Over and Cruz (all in this handbook), as well as to Hájek and Hall (1994) and Pfeifer and Douven (2014).

5. A critical discussion of another putative argument for probabilism, Cox's theorem, can be found in chapter 4.7 by Dubois and Prade (in this handbook).

6. There are unboundedly many references to both of these on the Internet. See also chapter 8.2 by Peterson (this handbook).

7. To be sure, the original Principal Principle holds as a constraint on rational agents more generally, assuming that they update by conditioning on evidence that is limited to the past and present—more on conditioning shortly. However, as we will see, the main arguments for conditioning have been questioned, as has conditioning itself. Lewis (1994) also offers a "New Principle" that explicitly supplants the original principle.

8. Carnap (1950) attempts to characterize a unique logical probability function and can thus also be seen as an early defender of Uniqueness. Various proponents of objective Bayesianism also fall into this camp (e.g., L. Williamson, 2010).

9. The literature on disagreement and higher-order evidence is vast. Matheson (2015) provides an overview of the debate. See also recent anthologies on the topic, such as Feldman and Warfield (2010) and Christensen and Lackey (2016).

References

Armendt, B. (1980). Is there a Dutch book argument for probability kinematics? *Philosophy of Science*, *47*(4), 583–588.

Arntzenius, F. (2003). Some problems for conditionalization and reflection. *Journal of Philosophy*, *100*(7), 356–370.

Bertrand, J. (1889). *Calcul des probabilités* [Calculus of probabilities]. Paris, France: Gauthier-Villars.

Briggs, R. (2009). Distorted reflection. *Philosophical Review*, *118*(1): 59–85.

Bronfman, A. (manuscript). *A gap in Joyce's argument for probabilism*. Unpublished manuscript.

Carnap, R. (1950). *Logical foundations of probability*. Chicago, IL: University of Chicago Press.

Christensen, D. (1991). Clever bookies and coherent beliefs. *Philosophical Review*, *100*(2), 229–247.

Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, *116*(2), 187–217.

Christensen, D., & Lackey, J. (Eds.). (2016). *The epistemology of disagreement: New essays*. Oxford, England: Oxford University Press.

De Bona, G., & Staffel, J. (2017). Graded incoherence for accuracy-firsters. *Philosophy of Science*, *84*(2), 189–213.

De Bona, G., & Staffel, J. (2018). Why be (approximately) coherent? *Analysis*, *78*(3), 405–415.

de Finetti, B. (1964). Foresight: Its logical laws, its subjective sources. In H. E. Kyburg & H. E. Smokler (Eds.), *Studies in subjective probability* (pp. 100–110). New York, NY: Wiley. (Original work published in French in 1937)

de Finetti, B. (1974). *Theory of probability*. New York, NY: Wiley.

Diaconis, P., & Zabell, S. (1982). Updating subjective probability. *Journal of the American Statistical Association*, *77*(380), 822–830.

Dogramaci, S., & Horowitz, S. (2016). An argument for uniqueness about evidential support. *Philosophical Issues*, *26*(1), 130–147.

Douven, I. (2009). Uniqueness revisited. *American Philosophical Quarterly*, *46*(4), 347–361.

Easwaran, K. (2013). Expected accuracy supports conditionalization—and conglomerability and reflection. *Philosophy of Science*, *80*(1), 119–142.

Easwaran, K. (2014). Regularity and infinitesimal credences. *Philosophical Review*, *123*(1), 1–41.

Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint*, *10*(5), 1–11.

Eriksson, L., & Hájek, A. (2007). What are degrees of belief? *Studia Logica*, *86*(2), 183–213.

Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (Ed.), *Philosophers without gods* (pp. 194–214). Oxford, England: Oxford University Press.

Feldman, R., & Warfield, T. A. (Eds.). (2010). *Disagreement*. Oxford, England: Oxford University Press.

Fishburn, P. C. (1986). The axioms of subjective probability. *Statistical Science*, *1*(3), 335–358.

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

Greco, D., & Hedden, B. (2016). Uniqueness and metaepistemology. *Journal of Philosophy*, *113*(8), 365–395.

Hájek, A. (2003). What conditional probability could not be. *Synthese*, *137*(3), 273–323.

Hájek, A. (2008). Arguments for—or against—probabilism? *British Journal for the Philosophy of Science*, *59*(4), 793–819.

Hájek, A. (2009). Dutch book arguments. In P. Anand, P. Pattanaik, & C. Puppe (Eds.), *The Oxford handbook of rational and social choice* (pp. 173–195). Oxford, England: Oxford University Press.

Hájek, A. (2012). Is strict coherence coherent? *Dialectica*, *66*(3), 411–424.

Hájek, A. (2019). Interpretations of probability. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Retrieved from https://plato.stanford.edu/archives/fall2019/entries/probability-interpret/

Hájek, A., & Hall, N. (1994). The hypothesis of the conditional construal of conditional probability. In E. Eells & B. Skyrms (Eds.), *Probability and conditionals* (pp. 75–111). Cambridge, England: Cambridge University Press.

Hájek, A., & Hitchcock, C. (2016). Probability for everyone—even philosophers. In A. Hájek & C. Hitchcock (Eds.), *The Oxford handbook of probability and philosophy* (pp. 5–30). Oxford, England: Oxford University Press.

Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge, England: Cambridge University Press.

Jeffrey, R. C. (1965). *The logic of decision*. Chicago, IL: University of Chicago Press.

Joyce, J. M. (1998). A non-pragmatic vindication of probabilism. *Philosophy of Science*, *65*(4), 575–603.

Joyce, J. M. (2005). How probabilities reflect evidence. *Philosophical Perspectives*, *19*, 153–178.

Kelly, T. (2005). The epistemic significance of disagreement. *Oxford Studies in Epistemology*, *1*, 167–196.

Kelly, T. (2013). Evidence can be permissive. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 298–311). Hoboken, NJ: Wiley.

Kemeny, J. G. (1955). Fair bets and inductive probabilities. *Journal of Symbolic Logic*, *20*(3), 263–273.

Kolmogorov, A. N. (1950). *Foundations of the theory of probability*. New York, NY: Chelsea. (Original work published in German as *Grundbegriffe der Wahrscheinlichkeitsrechnung* in 1933)

Kopec, M., & Titelbaum, M. (2016). The uniqueness thesis. *Philosophy Compass*, *11*(4), 189–200.

Lasonen-Aarnio, M. (2013). Disagreement and evidential attenuation. *Noûs*, *47*(4), 767–794.

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

Leitgeb, H., & Pettigrew, R. (2010b). An objective justification of Bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science*, *77*(2), 236–272.

Levi, I. (1967). Probability kinematics. *British Journal for the Philosophy of Science*, *18*(3), 297–315.

Levi, I. (1985). Imprecision and indeterminacy in probability judgment. *Philosophy of Science*, *52*(3), 390–409.

Levi, I. (1987). The demons of decision. *The Monist*, *70*(2), 193–211.

Levi, I. (1989). Possibility and probability. *Erkenntnis*, *31*(2–3), 365–386.

Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in inductive logic and probability* (Vol. II, pp. 83–132). Berkeley: University of California Press.

Lewis, D. (1994). Humean supervenience debugged. *Mind*, *103*(412), 473–490.

Lewis, D. (1999). Why conditionalize? In *Papers in metaphysics and epistemology* (pp. 403–407). Cambridge, England: Cambridge University Press.

Mahtani, A. (2015). Dutch books, coherence, and logical consistency. *Noûs*, *49*(3), 522–537.

Matheson, J. (2015). *The epistemic significance of disagreement*. New York, NY: Palgrave Macmillan.

Meacham, C. J. G. (2014). Impermissive Bayesianism. *Erkenntnis*, *79*(6)(Suppl.), 1185–1217.

Norton, J. D. (2008). Ignorance and indifference. *Philosophy of Science*, *75*(1), 45–68.

Novack, G. (2010). A defense of the principle of indifference. *Journal of Philosophical Logic*, *39*(6), 655–678.

Pettigrew, R. (2016). *Accuracy and the laws of chance*. Oxford, England: Oxford University Press.

Pettigrew, R. (2018, February 2). An almost-Dutch book argument for the Principal Principle. Retrieved from http://m-phi.blogspot.co.uk/2018/02/a-pragmatic-dutch-book-argument-for.html

Pfeifer, N., & Douven, I. (2014). Formal epistemology and the new paradigm psychology of reasoning. *Review of Philosophy and Psychology*, *5*(2), 199–221.

Popper, K. (1959). *The logic of scientific discovery*. New York, NY: Basic Books.

Ramsey, F. P. (1990). Truth and probability (original manuscript, 1926). In D. H. Mellor (Ed.), *Philosophical papers* (pp. 52–94). New York, NY: Cambridge University Press.

Rényi, A. (1970). *Foundations of probability*. San Francisco, CA: Holden-Day.

Rinard, S. (2015). A decision theory for imprecise probabilities. *Philosophers' Imprint*, *15*(7), 1–16.

Savage, L. J. (1954). *The foundations of statistics*. New York, NY: Wiley.

Schervish, M. J., Seidenfeld, T., & Kadane, J. B. (2002). Measuring incoherence. *Sankhya: The Indian Journal of Statistics*, *64*, 561–587.

Schoenfield, M. (2014). Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, *48*(1), 193–218.

Schoenfield, M. (2017). Conditionalization does not (in general) maximize expected accuracy. *Mind*, *51*(4), 667–685.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*(3), 379–423.

Shimony, A. (1955). Coherence and the axioms of confirmation. *Journal of Symbolic Logic*, *20*(1), 1–28.

Skyrms, B. (1980). Higher order degrees of belief. In D. H. Mellor (Ed.), *Prospects for pragmatism* (pp. 109–137). Cambridge, England: Cambridge University Press.

Skyrms, B. (1987). Dynamic coherence and probability kinematics. *Philosophy of Science*, *54*(1), 1–20.

Skyrms, B. (1993). A mistake in dynamic coherence arguments? *Philosophy of Science*, *60*(2), 320–328.

Slovic, P., & Lichtenstein, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, *6*(6), 649–744.

Stefánsson, H. O. (2017). What is "real" in probabilism? *Australasian Journal of Philosophy*, *95*(3), 573–587.

Talbott, W. (1991). Two principles of Bayesian epistemology. *Philosophical Studies*, *62*(2), 135–150.

Titelbaum, M. (2012). *Quitting certainties: A Bayesian framework modeling degrees of belief*. Oxford, England: Oxford University Press.

van Fraassen, B. C. (1981). A problem for relative information minimizers in probability kinematics. *British Journal for the Philosophy of Science*, *32*, 375–379.

van Fraassen, B. C. (1984). Belief and the will. *Journal of Philosophy*, *81*(5), 235–256.

van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford, England: Oxford University Press.

van Fraassen, B. C. (1999). Conditionalization, a new argument for. *Topoi*, *18*(2), 93–96.

Vineberg, S. (2001). The notion of consistency for partial belief. *Philosophical Studies*, *102*(3), 281–296.

Vineberg, S. (2016). Dutch book arguments. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Retrieved from https://plato.stanford.edu/archives/spr2016/entries/dutch-book/

White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, *19*, 445–459.

White, R. (2009). Evidential symmetry and mushy credence. In T. Szabó Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 161–186). Oxford, England: Oxford University Press.

Williamson, J. (2010). *In defense of objective Bayesianism*. Oxford, England: Oxford University Press.

Williamson, T. (2007). How probable is an infinite sequence of heads? *Analysis*, *67*(3), 173–180.

Zynda, L. (1996). Coherence as an ideal of rationality. *Synthese*, *109*(2), 175–216.

Zynda, L. (2000). Representation theorems and realism about degrees of belief. *Philosophy of Science*, *67*, 45–69.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Knauff, Markus, editor. | Spohn, Wolfgang, editor.
Title: The handbook of rationality / edited by Markus Knauff and Wolfgang Spohn.
Description: Cambridge : The MIT Press, 2021. | Includes bibliographical references and index.
Identifiers: LCCN 2020048455 | ISBN 9780262045070 (hardcover)
Subjects: LCSH: Reasoning (Psychology) | Reason. | Cognitive psychology. | Logic. | Philosophy of mind.
Classification: LCC BF442 .H36 2021 | DDC 153.4/3—dc23
LC record available at https://lccn.loc.gov/2020048455