Bayesianism in Science*

Josh Brekel† and Jonah N. Schupbach‡

Abstract

This chapter provides an introduction to Bayesianism and a sampling of its scientific applications. Section 1 summarizes the essential commitments of Bayesianism, including its interpretation of probability theory as a logic of consistency for degrees of belief. This section then clarifies some ways in which Bayesians diverge in their development and defense of further epistemological norms for rational degrees of belief. Section 2 surveys examples of the many ways in which Bayesianism is fruitfully applied in scientific practice, focusing on applications in cognitive science and cosmology.

Keywords: Bayesianism, Probability, Statistics, Confirmation Theory, Cognitive Psychology, Cosmology

Key Points:

- Bayesianism is a methodological school of thought which interprets probabilities as coherent degrees of belief or confidences.
- Bayesian epistemology explores and probabilistically explicates norms and concepts relevant to senses of rationality stretching beyond mere coherence.
- Bayesianism's interpretation of probability and Bayesian epistemological explications are scientifically fruitful tools, as revealed by examples from cognitive science.
- Bayesianism has also given rise to an immensely, scientifically useful school of statistics, as shown through examples from cosmology.

Introduction

This chapter provides an introduction to Bayesianism along with a discussion of some ways in which this methodological school of thought is fruitfully used by

*We owe special thanks both to Jeff Kasser and Tania Lombrozo for providing us with helpful feedback on earlier drafts of this chapter.

†Princeton University, Department of Psychology; brekel@princeton.edu.

‡University of Utah, Department of Philosophy; jonah.n.schupbach@utah.edu.

working scientists. Section 1 summarizes the essential commitments of Bayesianism, including its interpretation of probability theory as a logic of consistency for degrees of belief (§1.1). This section then clarifies some ways in which Bayesians diverge in their development and defense of further epistemological norms for rational degrees of belief (§1.2). Section 2 discusses examples of the many ways in which Bayesianism is applied in scientific practice. §2.1 focuses on applications of Bayesianism to cognitive science, while §2.2 offers examples from cosmology.

1 On Bayesianism

1.1 Probability Theory and the Bayesian Logic of Consistency

I. J. Good (1971) famously estimated there to be no less than 46,656 distinct varieties of Bayesianism, "more than the number of professional statisticians." It comes as no surprise then that it proves difficult to come up with a single, general definition of Bayesianism. Nonetheless, as a starting point, we may observe that Bayesianism is at least characterizable as a methodological school of thought involving a comparatively liberal application of *probabilities* to the study of uncertain inference. Bayesians are further unified in justifying their liberal use of probabilities by appeal to a particular philosophical interpretation of mathematical probabilities; for Bayesians, the probability calculus provides a logic of graded doxastic attitudes—alternately referred to as "degrees of belief," "credences," "confidences," and "uncertainties."

In adopting a probabilistic approach, Bayesians generally[1] accept Kolmogorov's (1933) canonical axiomatization of the probability calculus:

> Kolmogorov's Axioms.
> Where $\mathcal{A}$ is an algebra of propositions[2] and $\top \in \mathcal{A}$ is the tautology, $\Pr : \mathcal{A} \to \mathbb{R}$ is a probability function if and only if it satisfies the following three axioms:
>
> **Non-Negativity.** For any $\phi \in \mathcal{A}$, $\Pr(\phi) \geq 0$.
>
> **Normality.** $\Pr(\top) = 1$.
>
> **Finite Additivity.** For mutually exclusive $\phi, \psi \in \mathcal{A}$, $\Pr(\phi \vee \psi) = \Pr(\phi) + \Pr(\psi)$.

[1] "Generally" because various Bayesians take issue with virtually all of the mathematical features and axioms of Kolmogorov's approach. Lyon (2016) provides a helpful overview of such "discontents".

[2] $\mathcal{A}$ is an "algebra" of propositions iff $\mathcal{A}$ is a set of propositions having the tautology $\top$ as a member [$\top \in \mathcal{A}$], and being closed under classical logic's negation and disjunction operations [for any propositions ϕ and ψ, if $\phi, \psi \in \mathcal{A}$, then it is the case that $\neg\phi \in \mathcal{A}$ and $\phi \vee \psi \in \mathcal{A}$].

In cases for which $\mathcal{A}$ is a σ-algebra, closed under *countable* disjunctions,[3] Kolmogorov extends **Additivity** to the countably infinite case:

Countable Additivity. For mutually exclusive ϕ_i, $\Pr(\bigvee_{i=1}^{\infty} \phi_i) = \sum_{i=1}^{\infty} \Pr(\phi_i)$.

On top of these axioms, Kolmogorov adds the following *definition* of "conditional probability:"

Conditional Probability. For $\phi, \psi \in \mathcal{A}$, if $\Pr(\psi) > 0$, then:

$$\Pr(\phi|\psi) =_{def} \frac{\Pr(\phi \wedge \psi)}{\Pr(\psi)}.$$

A conditional probability, $\Pr(\phi|\psi)$, is meant to explicate the probability of ϕ "under the condition that" or "given" ψ.[4]

From these axioms and definition, all familiar rules and properties of probability are derivable, including the following:[5]

Negation. $\Pr(\neg\phi) = 1 - \Pr(\phi)$.

Contradiction. $\Pr(\bot) = 0$ for the logical contradiction $\bot \in \mathcal{A}$.

Maximality. For any $\phi \in \mathcal{A}$, $\Pr(\phi) \leq 1$.

Conjunction. $\Pr(\phi \wedge \psi) = \Pr(\phi)\Pr(\psi|\phi) = \Pr(\psi)\Pr(\phi|\psi)$.

General Additivity. $\Pr(\phi \vee \psi) = \Pr(\phi) + \Pr(\psi) - \Pr(\phi \wedge \psi)$.

[3] If $\phi_1, \phi_2, ... \in \mathcal{A}$, then $\bigvee_{i=1}^{\infty} \phi_i \in \mathcal{A}$.

[4] Kolmogorov's **Conditional Probability** is the source of much confusion and philosophical controversy. Here, we mention just a few important issues. First, many Bayesians maintain that **Conditional Probability** should *not* be thought of as formalizing the probability of the corresponding conditional, Pr(if ψ, then ϕ). As Lewis (1976) shows, identifying a conditional probability with the probability of a conditional (be it classical logic's material conditional, $\phi \rightarrow \psi$, or the English indicative conditional) trivializes the probability calculus. Second, Humphreys (1985) argues that **Conditional Probability** paradoxically rules out propensity interpretations of the probability calculus. Third, Hájek (2003) argues that this definition also conflicts with Bayesian interpretations of the calculus, given various "trouble spots," including the fact that genuine possibilities that we may meaningfully take as "given" can have measure zero. See (Bennett, 2003, ch. 5) and (Easwaran, 2019) for more discussion.

[5] For many of the derivations using this propositional-logical (as opposed to set-theoretic) framework, in addition to the axioms, one must invoke the assumption that logical equivalents are necessarily equiprobable. For example, here is a proof of **Contradiction**:

1. $\bot \dashv\vdash \neg\top$ (logical equivalence)
2. $\Pr(\bot) = \Pr(\neg\top)$ (1, equiprobability of logical equivalents assumption)
3. $\Pr(\top \vee \neg\top) = \Pr(\top) + \Pr(\neg\top) = 1$ (**Additivity**, **Normality**)
4. $\Pr(\neg\top) = 1 - \Pr(\top) = 0$ (3 algebra, **Normality**)
5. $\Pr(\bot) = 0$ (2,4 substitution) □

Total Probability. For n-wise partition[6] $\{\psi_1, \psi_2, ..., \psi_n\} \subseteq \mathcal{A}$, $\Pr(\phi) = \sum_{i=1}^{n} \Pr(\psi_i) \Pr(\phi|\psi_i)$.

Entailment. If $\phi \vdash \psi$, then $\Pr(\phi) \leq \Pr(\psi)$ and $\Pr(\psi|\phi) = 1$.

The mathematical centerpiece of Bayesianism, however, is the following theorem (presented here in both a simple and expanded form):

Bayes's Theorem (Simple Form). For $\phi, \psi \in \mathcal{A}$, if $\Pr(\psi) > 0$, then:

$$\Pr(\phi|\psi) = \frac{\Pr(\phi) \Pr(\psi|\phi)}{\Pr(\psi)}.$$

Bayes's Theorem (Expanded Form). For $\psi \in \mathcal{A}$, partition $\{\phi_1, \phi_2, ..., \phi_n\} \subseteq \mathcal{A}$, and $i \in \{1, 2, ...n\}$, if $\Pr(\psi) = \sum_{j=1}^{n} \Pr(\phi_j) \Pr(\psi|\phi_j) > 0$, then:

$$\Pr(\phi_i|\psi) = \frac{\Pr(\phi_i) \Pr(\psi|\phi_i)}{\sum_{j=1}^{n} \Pr(\phi_j) \Pr(\psi|\phi_j)}.$$

While Bayesianism is a controversial school of thought, there is nothing mathematically controversial about **Bayes's Theorem**. The reader can easily verify that the simple form of this theorem is straightforwardly derivable from our statement of **Conjunction** above, while the expanded form follows by using **Total Probability** to expand the simple form's denominator. Any probabilistic approach—be it Classical-Frequentist, Bayesian, Likelihoodist, or otherwise—is committed to **Bayes's Theorem** qua mathematical theorem simply by being committed to the probability calculus.

Probabilistic schools begin to part ways from each other when it comes to the question of how applicable **Bayes's Theorem** is in practice. It is in this sense that Bayesianism may be distinguished from its statistical rivals as involving a more liberal application of probabilities. As Edwards et al. (1963, p. 194) write in an early plea for psychologists to adopt Bayesian statistical methods more often in their work, "Bayesian statistics is so named for the rather inadequate reason that it has many more occasions to apply **Bayes's Theorem** than classical statistics has."

To clarify, note that a central goal of statistics and uncertain inference is to determine and test the bearing of evidence (or data) upon hypotheses. *Prima facie*, the obvious, salient probabilistic mark of such standing would simply be the relevant hypothesis's probability given the evidence, $\Pr(h|e)$.[7] **Bayes's Theorem**

[6] A partition is a set of mutually exclusive and jointly exhaustive propositions, such that exactly one member of the partition must be true.

[7] This standing is also, of course, determined in the light of "background knowledge" accepted as true. Accordingly, Bayesians explicate the salient probability as $\Pr(h|e \wedge b)$, where b stands for a propositional representation of that background information. That said, it is also common for Bayesians to leave the background term b implicit in all formulations. For the sake of notational simplicity and elegance, we follow this convention in this chapter.

provides the mathematical key for assessing this probability:

$$\Pr(h|e) = \frac{\Pr(h)\Pr(e|h)}{\Pr(e)}. \tag{1}$$

Bayesians agree that this theorem should and—perhaps more controversially—*can* indeed be a probabilistic focal point for hypothesis assessment. As a matter of terminology, Bayesians refer to $\Pr(h|e)$ as the hypothesis's "posterior probability," it being h's probability assessed in light of, or logically *posterior* to, evidence e. By contrast, $\Pr(h)$ is h's "prior probability." Across Bayesian and non-Bayesian camps, $\Pr(e|h)$ is termed h's "likelihood" with respect to e, or more fully, the "likelihood of e given h."

As Gelman and Shalizi (2013, p. 9) summarize when describing the "usual" view of Bayesianism: "The expression [$\Pr(h|e)$] says it all, and the central goal of Bayesian inference is computing the posterior probabilities of hypotheses." Examining the right-hand side of (1), however, it is less than clear how to evaluate such probabilities, or whether some of these probabilities are meaningful at all. Consider the probability of the evidence obtaining, $\Pr(e)$. In scientific applications such as those we review in later sections, the evidence is often qualitative, imprecise, singular, or unrepeatable in ways that seem to make assessments of its probability complete guesswork at best and meaningless at worst.

One may attempt to bypass this problem by using the expanded form of **Bayes's Theorem**:

$$\Pr(h_i|e) = \frac{\Pr(h_i)\Pr(e|h_i)}{\sum_{j=1}^{n}\Pr(h_j)\Pr(e|h_j)}. \tag{2}$$

This expanded form makes no mention of $\Pr(e)$ when applied to the case of assessing a hypothesis's probability in light of evidence. However, the same challenge of enigmatic probabilities rears its head with respect to the so-called *prior probabilities* of hypotheses, $\Pr(h_j)$, occurring throughout the right-hand side of (2). Is there a general, meaningful sense in which we may talk about the probability of a hypothesis, logically prior to consideration of the evidence?

So-called "physical" interpretations—e.g., relative frequency (von Mises, 1957) or propensity interpretations (Popper, 1959b)—understand probabilities to be "objective properties" of physical objects, arrangements, or events in "the real world" (Wasserman, 2004, p. 175). The relative frequency interpretation commonly accepted by Bayesianism's Frequentist rivals, for instance, interprets an event's probability to be the limiting, long-run relative frequency of that event's occurrence within a more general "reference class" of events. Such interpretations render talk of a hypothesis's probability meaningless, thereby implying a negative response to the above question. For one thing, hypotheses don't occur or not; rather, they are true or not. For another thing, even if we allow *a hypothesis being true* into the ranks of occur-able events, there is rarely a way to make sense of the proportion of times that a hypothesis (often a general claim on nature) has been true in the

long-run over some repeatable reference class of events or universes. As Peirce famously makes the point from his Frequentist perspective, "The relative probability of this or that arrangement of Nature is something which we should have a right to talk about if universes were as plenty as blackberries" (Peirce, 1992, p. 165).

Bayesians, by contrast, are unified by their affirmative answer to the above question and consequent willingness to use full probability distributions, assigning values to prior probabilities of hypotheses. As Edwards et al. (1963, p. 195) write, classical statistical approaches "made useless, if not meaningless, the probability that a hypothesis is true, [and thus were forced to develop] techniques for statistical inference which did not attach probabilities to hypotheses." Bayesianism instead is committed to a "reinstatement of the probabilities of hypotheses"—which Edwards et al. argue effects "a return of statistical inference to its original line of development." Similarly, Good (1971) suggests that a bond unifying his 46,656 varieties of Bayesianism is the belief "that it is usually meaningful to talk about the probability of a hypothesis." Bayesian methods thus have the distinguishing feature of employing "full probability models—joint probability distributions for all observable and unobservable quantities in a problem" (Gelman et al., 2014, p. 3).

Bayesians justify their appeal to prior probabilities of hypotheses and applications of **Bayes's Theorem** by way of their radically different interpretation of mathematical probabilities. Bayesians understand probabilities to be *psyche*-logical or epistemological instead of physical.[8] Here, probabilities are interpreted as normative counterparts to actual psychological phenomena. "Probability is *orderly* opinion" (Edwards et al., 1963, p. 194) or "coherent opinion" (de Finetti, 1937, p. 104). More specifically, probabilities are interpreted as *consistent* or *coherent* "degrees of belief," "credences," "uncertainties," or "confidences."[9] On this interpretation, insofar as it is meaningful to talk about how confident one is that various propositions are true, we may inquire whether such confidences are mutually coherent. The Bayesian's suggestion is that probabilities can usefully be interpreted as such coherent degrees of confidence. The Bayesian interpretation of probability ostensibly renders meaningful the probability of a hypothesis (as well as the probability that the evidence attains, along with the probability of any other meaningful proposition); $\Pr(h)$ explicates a degree of belief or confidence that an agent may coherently have in h, prior to e, but in light of the background knowledge accepted by that agent.

[8]While an overview of the various defenses that Bayesians offer for this interpretation goes beyond the scope of this chapter, there is an overabundance of Bayesian primers that interested readers might consult for more information; including any of the following: (Howson and Urbach, 2006; Easwaran, 2011; Schupbach, 2022; Titelbaum, 2022; Lin, 2024).

[9]Bayesians often alternate between these labels, even sometimes within a single work (Ramsey, 1926; de Finetti, 1937; Savage, 1954). In this chapter, we allow ourselves a similar leniency in the terminology we use to refer to the graded doxastic states explicated by probabilities.

On this interpretation, the probability calculus itself provides a "logic of consistency" for degrees of belief had by an agent synchronically, or at a single time (Ramsey, 1926). For example, Kolmogorov's **Normality** axiom is interpreted as requiring that an agent ought to have maximal confidence (i.e., certainty) in any logical truth $\top$; **Negation** builds off of this norm by requiring that an agent's confidences in any proposition ϕ and its negation $\neg\phi$, whatever they be, must sum to one; and so on. Indeed, probability theory can be understood as a conservative, straightforward generalization of classical, deductive logic's bivalent, truth-functional semantics for cases in which, instead of truth-values, we are dealing with continuous degrees of uncertainty or confidence (Cox, 1961; Jaynes, 2003).

1.2 Bayesian Epistemology

While the probability calculus constrains degrees of belief on pain of inconsistency, many Bayesians go further in developing and proposing probabilistic norms indicative of a sense of rationality stretching beyond mere consistency. Bayesian epistemology (Bovens and Hartmann, 2003) concerns further normative questions regarding which coherent graded belief states an agent ought to occupy and how those states ought to evolve in response to evidence. As Putnam (1981, pp. 192–3) puts it, the aim is to flesh out a "formalization of the psychology of an ideally rational human scientist."

The most famous such norm, which was suggested by Thomas Bayes (1763) in his seminal essay (Earman 1992, ch. 1; Schupbach 2022, §2.2), is so engrained in the traditional Bayesian approach that it is often taken to be an additional defining commitment of Bayesianism. Now known as "Bayes's Rule" or "Bayesian Conditionalization," this norm proposes a *diachronic* notion of coherence, legislating how a rational agent should consistently update their uncertainties across time with the intake of new evidence:

Bayes's Rule. Upon learning ψ (and nothing else), the new probability of any proposition $\phi \in \mathcal{A}$ should equal the old (just previous to learning ψ) conditional probability of ϕ given ψ: $\Pr_{new}(\phi) = \Pr_{old}(\phi|\psi)$.

Bayes's Rule should not be confused with **Bayes's Theorem**; the former is no consequence of the the probability axioms, nor is it even a constraint within a probability distribution. Rather, **Bayes's Rule** relates two distinct probability distributions to one another. In terms of a particular hypothesis h, this norm mandates that a rational agent who learns e (and nothing else) across two times will update her confidence in h, $\Pr_{new}(h)$, to equal precisely the same degree as her prior confidence in h (prior to taking in that new evidence) conditional on e, $\Pr_{old}(h|e)$.[10] The rational agent's confidences in *all* other propositions are similarly updated

[10]Note that the Bayesian terminology of "prior" and "posterior probabilities" takes on additional, temporal significance in light of **Bayes's Rule**.

in light of e, resulting in a new probability distribution $Pr_{new}(\cdot)$ representing her rational confidences after learning e.

By itself, **Bayes's Rule** places increasingly demanding constraints on rational credences as more and more evidence comes in. Indeed, some Bayesians have argued that this norm suffices to objectify or "swamp" altogether the initial subjectivism permitted by Bayesianism's logic of coherence. The idea is expressed clearly in various "convergence theorems" (Blackwell and Dubins, 1962; Gaifman and Snir, 1982), which show that intersubjective differences in confidences between rational agents diminish (*entirely* in the long-run) as the agents all update on the same evidence using **Bayes's Rule**. This epistemological interpretation of such theorems has, however, been called into question; Schupbach (2022, pp. 47-49), for example, argues that other epistemological norms (which calibrate likelihoods) really do the heavy-lifting in such results. For more concerns and discussion surrounding the significance of Bayesian convergence theorems, the reader is directed to (Earman, 1992, ch. 6) and (Huttegger, 2015).

While **Bayes's Rule** legislates the updating of degrees of belief over time, other principles constrain rational degrees of belief at a time. The most venerable such *synchronic* principle—tracing back to Laplace (1814) and made more famous by Keynes (1921)—is arguably the principle of **Indifference** (also known as "Insufficient Reason" and generalized by such principles as "Maximum Entropy"):

Indifference. If there is no known reason favoring any one member of a set of mutually exclusive propositions over the others,[11] then we should assign equal probability to each member.

In the absence of evidence favoring any particular outcome, this principle requires an agent to have equal confidence that a normal-looking six-sided die, for example, will land on any particular side. It thus constrains rational degrees of belief to precise, equal values in cases where an agent's evidence is neutral.

This principle evidently seems epistemologically secure, insofar as any other degrees of belief to such options would appear to be "an expression of sheer prejudice" (Jeffreys, 1939, p. 33). Nonetheless, as with **Bayes's Rule**, **Indifference** is the subject of much controversy. Most famously, Bertrand's (1889) Paradox shows that this principle leads to inconsistent, equally defensible assignments of probabilities to the same propositions. See (van Fraassen, 1989, p. 303) for a more recent, influential presentation of this paradox, and (Williamson 2010, ch. 9; Schupbach 2022, §2.4) for in-depth explorations of these issues including potential responses in defense of **Indifference**.

While a complete survey of further, proposed norms of Bayesian epistemology falls outside the scope of this chapter, such norms may generally be categorized along the following lines:

[11] Or, in other versions, if the total evidence no more supports any one of the propositions over the others.

- Principles constraining prior probabilities (**Regularity**, **Indifference**, **Maximum Entropy**, etc.),
- Principles legislating how a rational agent ought to update their probabilities diachronically with the intake of new evidence (**Bayes's Rule**, **Jeffrey's Rule**, Titelbaum's (2013) **Generalized Conditionalization**, etc.), and
- Principles regarding how a rational agent's confidences ought to defer to known objective chances, statistical frequencies, and expert sources (**The Principal Principle**, **Direct Inference**, **Reflection**, etc.).

To the extent that the Bayesian accepts such epistemological principles, rational degrees of belief are constrained beyond the norms of coherence provided by the probability theory itself. This observation implies a spectrum of Bayesian views (Weisberg, 2009, §3). On the far left (liberal) side of the spectrum, "subjective Bayesians" merely require coherence for rationality; so long as an agent's confidences are jointly explicable within a single probability distribution, they are rational (de Finetti, 1937; Savage, 1954; Jeffrey, 1983).[12] This position is maximally permissive insofar as it permits any amount of intersubjective variation between probabilistically coherent, intrasubjective confidences. On the far right side of the spectrum, "objective Bayesians" defend some combination of principles that collectively, fully constrain rational degrees of belief (Jeffreys, 1939; Jaynes, 1957; Carnap, 1962; Rosenkrantz, 1981; Williamson, 2010). That is, on such a position, agents' confidences are forced upon them by their background knowledge, evidence, and the constraints of rationality; such confidences are only rational if they are jointly explicable within *the* uniquely appropriate probability distribution (the one distribution satisfying that maximally constraining set of principles).[13]

Bayesian epistemology is also characterized by the application of Bayesian probabilities to the study of epistemological concepts and questions. To take one example, "Bayesian confirmation theory" probabilistically explicates the concept of evidential support or "incremental confirmation."[14]

[12]Note that this *does not* imply that an agent's actual confidences must be real-value precise in order to be rational. Imprecise confidences (be they qualitative, comparative, etc.) or quantifiably precise confidences may equally count as rational so long as they are jointly explicable within a single probability distribution (Schupbach, 2022, §1.4.2).

[13]Echoing footnote 12 then, imprecise confidences (be they qualitative, comparative, etc.) or quantifiably precise confidences may equally count as rational for the objective Bayesian too, so long as they are jointly explicable within *the* uniquely appropriate probability distribution.

[14]"Confirmation" is actually understood to be polysemous in the confirmation theory literature. Going back to (Carnap, 1962, p. xvi), Bayesians distinguish the notion of incremental confirmation from the following *absolute* sense of confirmation—explicating h being "proved" or confirmed beyond a reasonable doubt in light of e:

Absolute Confirmation. h is confirmed if and only if $\Pr(h|e) > t$, where t is a suitable threshold signifying a value sufficiently high to put h beyond reasonable dispute.

Several Bayesian epistemologists have also helpfully noted the distinction between absolute confir-

Incremental Confirmation. Hypothesis h is confirmed by e (to some extent) if and only if e increases h's probability, $\Pr(h|e) > \Pr(h)$.

This formal account naturally suggests corresponding explications of disconfirmation and "confirmational irrelevance:"

Incremental Disconfirmation. Hypothesis h is *dis*confirmed by e (to some extent) if and only if e decreases h's probability, $\Pr(h|e) < \Pr(h)$.

Confirmational Irrelevance. Evidence e is confirmationally irrelevant to h if and only if $\Pr(h|e) = \Pr(h)$.

These qualitative, Bayesian accounts naturally suggest that one might measure *the degree to which* some e incrementally confirms an h by gauging the strength of the inequality between $\Pr(h|e)$ and $\Pr(h)$. And indeed, a multitude of distinct (i.e., ordinally inequivalent) Bayesian measures of degree of confirmation have been proposed and defended in the last many decades tracing back to foundational work by early Bayesians (Keynes, 1921; Kemeny and Oppenheim, 1952; Carnap, 1962) up through more recent times (Fitelson, 1999, 2001; Crupi et al., 2007).

§2 provides examples of the scientific import of Bayesianism in its various guises outlined above. That is, the unifying Bayesian idea that probabilities might be used to formalize a notion of orderly, coherent attitudes of uncertainty has fruitful application across sciences. However, as we will show by example, so does the Bayesian epistemological approach combining this interpretation of probability with more stringent epistemological principles of rationality. And indeed, some Bayesian accounts of epistemological concepts have been found to have fruitful scientific application. For example, as we summarize in §2.1.1, the Bayesian explication of incremental confirmation has been used by cognitive psychologists in their study of reasoning fallacies.

2 Bayesianism in Science

An exhaustive survey of the scientific uses of Bayesianism is not feasible—and not merely for a chapter of this length. In fact, Bayesianism has laid the theoretical foundations for much of the current work across certain social sciences. For example, the work of prominent Bayesians (von Neumann and Morgenstern, 1944; Savage, 1954; Luce and Raiffa, 1957) developing the formal framework of decision theory, including expected utility theory and game theory, grounds modern microeconomics. With that framework in place, economists are better able to model and investigate features of the economy. Economic game theory has now spilled

mation (as explicated by a sufficiently high posterior probability) and "weight of evidence," roughly corresponding to the amount of evidence backing a particular probability assignment (Keynes, 1921; Joyce, 2005; Kasser, 2016; Brekel, 2025).

over into the natural sciences (through evolutionary game theory), revealing one way in which Bayesianism extends its reach beyond the social sciences. This section will eventually discuss Bayesianism's now-commonplace application in the natural sciences. But we begin by discussing examples from one area of the social sciences.

Cognitive science provides a particularly rich case study for considering scientific uses of Bayesianism. Cognitive science pursues the interdisciplinary study of the mind, bridging social sciences like psychology with so-called hard sciences like neuroscience. Unlike many domains—where Bayesian statistics operates as a tool for data analysis—cognitive science treats the Bayesian interpretation of probability as a mainline hypothesis to be investigated. Oaksford and Chater (2007, p. 80)—two prominent voices in the "new paradigm" of reasoning—make this point emphatically:

> The subjective [Bayesian] interpretation of probability is essential, not just for the application of probability to scientific inference, but to the application of probability to *cognition*. In typical applications, as we shall see below, the cognitive system is viewed as weighing up evidence, e.g. from perception or memory, concerning whether or not some hypothesis, categorization, or causal relation holds. This can be understood in terms of degree of belief—but not, of course, in terms of limiting frequencies, or objective properties of the external world.

Bayesianism's role in the study of the mind differs across cognitive science. Here, we consider two subareas of cognitive science. First, Bayesianism plays a defining role in the study of judgment and decision-making. Second, Bayesian methods provide the framework underlying the so-called "new paradigm" in cognitive psychology, which reconceives human reasoning as fundamentally probabilistic (rather than deductive, for example).

While we choose to focus on these two subareas of cognitive science, Bayesianism's reach extends further through the discipline. For example, Bayesianism fuels computational models of the mind.[15] Pitched at Marr's "computational" level of analysis (Marr, 2010, ch. 1), these models offer Bayesian accounts of wide-ranging cognitive processes such as visual perception, animal cognition, concept learning, human and artificial (i.e., machine) understanding of causality, and much more. See (Griffiths et al., 2023) for a helpful introduction to this area of cognitive science.

[15] Another area we do not discuss, but which exemplifies Bayesianism's potential application to the neuroscientific study of the mind, includes work on the "Bayesian brain." In this field, neuroscientists explore the theory that the brain functions as a Bayesian, "predictive machine." Hohwy (2014), Mangalam (2025), and Bottemanne (2025) provide helpful reviews and discussions of such work.

2.1 Bayesian Epistemology in Cognitive Science

2.1.1 Judgment and Decision-Making

Any introduction to Bayesian cognitive science must at least mention the groundbreaking work of psychologists Amos Tversky and Daniel Kahneman. Over the course of 50 years, Tversky and Kahneman's "heuristics and biases research program" (HBP) has revolutionized the ways in which psychologists, epistemologists, economists, business leaders, and policy-makers think about human judgment and decision-making. Here, we provide an introductory gloss on HBP; more thorough reviews are offered by Tversky and Kahneman (1974), Kahneman et al. (1982), Gilovich et al. (2002), Manktelow (2012), and—for popular press treatments—Kahneman (2011), Thaler and Sunstein (2008), Kahneman et al. (2022), or Lewis (2017).

The fundamental motivation of HBP was to establish the existence of what Stanovich (1999) calls the *normative-descriptive gap*, which is "the gap between the prescriptions of normative theories and tested human performance"(Manktelow, 2012, p. 254). Consequently, Kahneman and Tversky designed experiments exploring how actual human reasoners fare in comparison to ideal Bayesian agents, particularly with respect to rationality, decision-making, and problems of prediction. In this way, HBP takes Bayesian epistemology as a starting point: an ideal benchmark against which they compare actual human judgments under uncertainty. After finding (what they argue are) widespread patterns of deviation from the Bayesian ideal, Tversky and Kahneman used this research to conclude that human reasoning diverges from Bayesian rationality in predictable ways.

The best way to understand the HBP is by considering an example of the normative-descriptive gap. The HBP provides many such examples from which to choose. Here, we focus on "the conjunction fallacy," as it illustrates a paradigmatic case of HBP research. In experiments regarding the conjunction fallacy, Tversky and Kahneman provided their participants with background information that affected the subsequent judgments of the participants. In arguably the most famous example, this information included the following personality sketch:

> Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations (Tversky and Kahneman, 1983, p. 297).

In one variation of their experiments, after reading this, participants were asked to rank associated statements about Linda in terms of their respective probabilities. So, for example, respondents were asked to consider whether it would be more probable for Linda to be a member of the League of Women Voters or for Linda to be an insurance salesperson.

Three such statements were of the most interest to Kahneman and Tversky:

F**:** Linda is active in the feminist movement.

B**:** Linda is a bank teller.

$B \wedge F$**:** Linda is a bank teller and is active in the feminist movement.

According to the probabilistic theorem of **Conjunction** (introduced above in §1.1),

$$\Pr(B \wedge F) = \Pr(B) \times \Pr(F|B) = \Pr(F) \times \Pr(B|F).$$

Since $\Pr(F|B)$ and $\Pr(B|F)$ must each fall between 0 and 1 (see **Non-Negativity** and **Maximality** from §1.1), **Conjunction** implies that "a conjunction cannot be more probable than one of its constituents" (Tversky and Kahneman, 1983, p. 294). Indeed, **Conjunction** implies that $\Pr(B \wedge F) \leq \Pr(B)$ and likewise that $\Pr(B \wedge F) \leq \Pr(F)$.[16] Intuitively, this makes sense since $B \wedge F$ is only true for cases in which B and F are both individually true, while B and F have broader truth-conditions, each possibly being true in cases where the other is false.

Contrary to Bayesian norms of coherence, Tversky and Kahneman found that respondents routinely violated **Conjunction**; for the "Linda problem" in particular, respondents tended to rank $B \wedge F$ as more probable than B. Tversky and Kahneman report that even among "sophisticated" subjects—e.g., seasoned doctoral students in the decision science program at the Stanford Business School—**Conjunction** was violated in this way no less than 85% of the time. These results become more shocking once it is recognized that such subjects had taken multiple classes on probability theory. Building upon their main study, Tversky and Kahneman found that even after "a series of increasingly desperate manipulations designed to induce subjects to obey the conjunction rule [...] 85% of respondents indicated that [$B \wedge F$] was more probable than [B], in a flagrant violation of the conjunction rule" (Tversky and Kahneman, 1983, p. 299). As a result, they concluded that descriptive analyses in psychology, economics, or other social sciences that assume adherence to the laws of probability provide inadequate descriptions of human judgment (Tversky and Kahneman, 1983, p. 313).

So far, our discussion of the Linda case has focused on the negative agenda (i.e., the "bias" portion) of the heuristics and biases program. While the word "bias" is used to describe the systematic departures from normative theory purportedly demonstrated in the research program, the word "heuristics" refers to the hypothesized explanations for those systematic errors. As Manktelow explains, "A heuristic is a reliable but not totally precise decision rule, a rule of thumb" (Manktelow, 2012, p. 54). Rather than reasoning probabilistically, then, the positive agenda of HBP asserts that humans make judgments on the basis of rough rules that are categorically different from the rules specified by Bayesian models of rationality (Gilovich et al., 2002, p. 3–4). Nonetheless, Bayesianism remains central to research on heuristics, insofar as it sets the normative standards such heuristics are meant to approximate.

[16] Note that the **Entailment** theorem listed in §1.1 implies this directly, since $B \wedge F \vdash B$ and $B \wedge F \vdash F$.

In the case of the conjunction fallacy, Tversky and Kahneman hypothesized that respondents used a "representativeness heuristic" to guide their assessments of probability. Put simply: Linda's personality profile is more representative of the proposition that she is a feminist bank teller in comparison to the less informative conjecture that Linda is merely a bank teller. If humans judge the probability of an event "by the degree to which that event is representative of an appropriate mental model," (Tversky and Kahneman, 1983, p. 295), then they use the representativeness heuristic when evaluating probability claims. Although work in HBP later generalized beyond the representativeness heuristic,[17] the interplay between the representativeness heuristic and the conjunction fallacy provides a canonical example of the sort of research undertaken in HBP.

It is difficult to overstate the influence of Kahneman and Tversky's HBP. For a few examples, HBP-influenced research spans from (Nisbett, 1980) in social psychology to (Jolls et al., 1998) in legal theory and from (Thaler, 1980) in economics to (Redelmeier et al., 1993; Crupi et al., 2018) in clinical medicine. Of course, with great fame comes great notoriety; many criticisms and developments of HBP arose in response to Tversky and Kahneman's research.[18] It is important to note, however, that these critical responses to the HBP are themselves often motivated by Bayesianism.

For example, Gigerenzer (1991) famously argues that people are naturally much closer to coherent probabilistic reasoners than suggested by many of the HBP's results, including those with respect to the conjunction fallacy. Participant judgments depart from the Bayesian norms in conjunction tasks not because of a systematic fallacy in their reasoning but because these tasks are framed using the unfamiliar, "ecologically invalid" format of percentages and numerical probabilities. Gigerenzer and Hoffrage (1995, p. 700) point out that when these same tasks are instead framed using "natural frequencies", "the 'conjunction fallacy' in the Linda problem and similar problems largely disappear[s]" (Fiedler, 1988; Gigerenzer, 1991; Gigerenzer and Hoffrage, 1995). Thus, a careful consideration of how to present information to people in an ecologically valid format can "help people reason the Bayesian way without even teaching them" (Gigerenzer and Hoffrage, 1995, p. 685). The upshot of such research is arguably that the normative-descriptive gap between Bayesian norms and human judgments is much narrower than the HBP often suggests. Of course, proponents of HBP have offered various responses to Gigerenzer and colleagues (Kahneman and Tversky, 1996; Gilovich et al., 2002; Manktelow, 2012).

Another account of the conjunction fallacy is worth highlighting as it builds directly off of work in Bayesian epistemology. Crupi, Fitelson, and Tentori (2008) consider whether Bayesian confirmation theory—as briefly summarized in §1.2—may

[17] Gilovich et al. (2002) develop and present a general system of heuristics used in reasoning.

[18] For taxonomies of critical discussions of Kahneman and Tversky's work, see (Stanovich, 1999), (Gilovich et al., 2002), and (Manktelow, 2012).

go some way toward explaining the reasoning patterns observed in classic instances of the conjunction fallacy. Informally, the idea behind their account is that experimental participants who fallaciously judge $B \wedge F$ to be "more probable" than B are accurately tracking confirmation relations.[19] In such cases, it is intuitive that the information respondents read profiling Linda's personality better confirms the conjecture that she is a feminist bank teller ($B \wedge F$) than it does the weaker conjecture that she is a bank teller (B).

To back up such intuitive, informal claims, Crupi, Fitelson, and Tentori observe that conjunction fallacy cases generally share three conditions, each straightforwardly explicable in Bayesian confirmation-theoretic terms. Regarding the Linda example specifically, where L stands for Linda's personality profile, these conditions and their confirmation-theoretic explications are as follows:

(i) L is negatively (if at all) correlated with B; i.e., L either disconfirms or is confirmationally irrelevant to B: $\Pr(B|L) \leq \Pr(B)$.

(ii) L is positively correlated with F, given B; i.e., L confirms F in light of B: $\Pr(F|L \wedge B) > \Pr(F|B)$.

(iii) B and F are mildly (if at all) negatively correlated; i.e., F either mildly disconfirms or is confirmationally irrelevant to B:: $\Pr(B|F) \leq \Pr(B)$.

As Crupi, Fitelson, and Tentori prove, conditions (i) and (ii) probabilistically imply the result that $B \wedge F$ is better confirmed by L than is B; where $\mathcal{C}$ may stand for any of the candidate Bayesian measures of degree of incremental confirmation: $\mathcal{C}(B \wedge F, L) > \mathcal{C}(B, L)$. Consequently, they hypothesize that participants in conjunction studies may accurately track such confirmation relations when they offer their (still incorrect) probability judgments. As Crupi, Fitelson, and Tentori conclude: "the hypothesis is that, on conditions (i) (iii), most participants may depart from the relevant probabilistic relationship [in the Linda case, between $\Pr(B \wedge F|L)$ and $\Pr(B|L)$] because of the perception that [$\mathcal{C}(B \wedge F, L) > \mathcal{C}(B, L)$]."

This Bayesian confirmation-theoretic account of the conjunction fallacy has proven itself robust in the face of criticism (Schupbach, 2012; Tentori and Crupi, 2012) and superior to some other accounts of the conjunction fallacy (Tentori et al., 2013). By drawing on the tools of Bayesian confirmation theory, Crupi, Fitelson, and Tentori's research effectively leverages a different normative concept (confirmation) via which one may assess reasoning. In terms of purely probabilistic reasoning, the participants in the HBP experiments (arguably) reasoned fallaciously. That said, their responses may be deemed rational with respect to the standards

[19] In a similar vein, Hintikka (2004, p. 32) argues that "given suitable collateral evidence, certain messages are self-confirming in that they tend to increase our confidence in their source." Accordingly, Hintikka (2004) claims that the reasoning at work in the conjunction experiments is not fallacious—even given Bayesian norms—because the propositions in question justify the use of distinct probability distributions. See Levi (2004, p. 37–8) for criticisms of this proposed solution.

specified by Bayesian confirmation theory.[20] This general approach illuminates other possible ways to use Bayesian explications to vindicate lay reasoning. Recently, for example, Jönsson and Shogenji (2019) appeal to Bayesian explications of the concept of coherence (Shogenji, 1999; Bovens and Hartmann, 2002; Schippers and Siebel, 2021) to argue that a Bayesian *coherence*-theoretic account of the conjunction fallacy improves upon Crupi, Fitelson, and Tentori's account. Plausibly, recent Bayesian work on "conjunctive explanations" (Schupbach and Glass, 2023) and explanatory goodness (Good, 1968; Glass, 2023; Schupbach, 2023) could likewise advance our understanding of this important fallacy.

2.1.2 "The New Paradigm"

By contrast with the HBP's attempt to explore the breadth of a postulated gap between Bayesian norms and human reasoning, a movement in psychology (beginning in earnest in the 1990s) coalesced around the idea that Bayesianism might provide fruitful *descriptive* models of human cognition (Elqayam and Over, 2013). Cognitive scientists working in this "new paradigm" placed Bayesianism's graded (non-binary) doxastic attitudes at the center of their approach to understanding human reasoning (Oaksford and Chater, 2009; Elqayam and Over, 2013; Douven, 2022).

Mention of a "new paradigm" raises questions concerning the old paradigm. The "old paradigm," in this instance was grounded in the idea that the norms of rationality were given by classical, deductive logic (Douven, 2022). On these foundations, "Descriptive Deductivism" worked on the assumption that, if humans regularly reason rationally, then we should observe them adhering to the norms of deductive logic. Humans reason as if their aim was to obey the norms of deductive logic, and thus the tools of deductive logic may provide fruitful explications of human reasoning.

Talk of a paradigm shift here can be illuminated by Kuhn (2012)'s seminal work on the topic. According to Kuhn (2012), shifts between paradigms occur when a critical mass of anomalies (e.g., observations that the existing paradigm struggles to account for) accumulates. New paradigm cognitive scientists argue that Descriptive Deductivism fell prey to just such a mass of anomalies, which were then convincingly resolved within the Bayesian paradigm. We here limit

[20] Arguably, there is still a sense in which the participants made a mistake, even after accounting for this vindication of their reasoning: when asked to provide answers in terms of *probabilities*, the participants gave their answers in terms of *confirmation theory*. However, there is some room for philosophical debate here. For example, Cohen (1979, 1981) famously argues that these types of results impugn on the formal explication of probability rather than on lay reasoning. Furthermore, (Levi, 2004, p. 38) argues that for the conjunction effect to be deemed fallacious, "We would have to be sure that the experimental subjects are evaluating probabilities obeying the calculus of probabilities and not some other index of support on the evidence. They are asked to *rank* the propositions [rather than being asked to evaluate the propositions using the probability calculus]." Perhaps the participants do not rank the propositions in terms of pure probabilities.

our focus to one such anomaly, arising with respect to the Wason Selection Task (WST). Not only is the WST "the most discussed task in philosophical debates about rationality", but it is also "a laboratory version of the problem of choosing the best experiments to test scientific laws"(Oaksford and Chater, 2007, 166).[21]

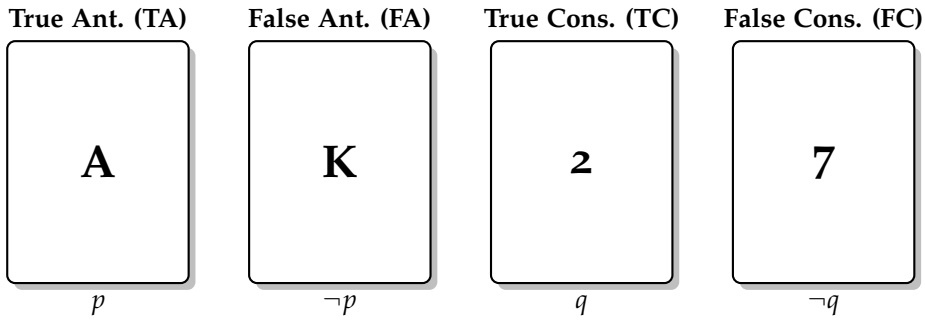


Figure 1: Wason Selection Task with standard labels for each card. The typical rule to test is: *If a card has a vowel on one side, it has an even number on the other* ($p \rightarrow q$) (Wason, 1968, p. 273). According to the old paradigm, cards should be selected to falsify the rule (i.e., select cards **A** and **7**).

In the classic WST study, Wason presented subjects with four cards—with a letter on one side and a number on the other—only one side of each being visible. Participants were asked which cards they would need to check (i.e., flip) in order to test the truth of the following general rule: *If a card has a vowel on one side, it has an even number on the other* (Wason, 1968, p. 273). As Wason himself put it, "The task was to select all those cards, but only those cards, which would have to be turned over in order to discover whether the experimenter was lying in making the conditional sentence" (Wason, 1968, pp. 273).

Using classical deductive logic, the rule is explicated using the material conditional, $p \rightarrow q$, defined truth-functionally in the following truth table:

p	q	$p \rightarrow q$
True	True	True
True	False	False
False	True	True
False	False	True

This table indicates which cards one would need to select to show that the rule is false, $p \rightarrow q$ being false if and only if the antecedent p is true and the consequent q is false. According to deductive logic, then, the best approach would be to attempt to falsify the rule by turning card p showing **A** and card $\neg q$ showing **7**;

[21] The choice of one example of work in the new paradigm is difficult, in part because "contributions to the new paradigm are rich and varied, and it is hard to give a brief summary about it that everyone contributing to it would accept" (Elqayam and Over, 2013, p. 258). We chose to focus on the WST due to the close connection between it and the philosophy of scientific methodology.

indeed, for the sake of falsifying the material conditional, it is entirely irrelevant what the alternate faces of cards $\neg p$ and q would reveal.[22] That is not what happened. As Wason explains, "Nearly all subjects select [p], from 60 to 75% select [q], only a minority select [$\neg q$] and hardly any select [$\neg p$]. Thus two errors are committed: the consequent is fallaciously affirmed and the contrapositive is withheld" (Wason, 1968, pp. 273–4). Unless participants are reasoning hopelessly fallaciously then, these results constitute a significant challenge to Descriptive Deductivism.

In the wake of repeated failed attempts to explain these results in the Descriptive Deductivist framework (Douven, 2022, p. 71), Oaksford and Chater (1994) provided a novel account of the WST—this account widely regarded as the opening salvo of the new paradigm. They judged their project to be one of rational analysis. Rational analysis stems from the work of Anderson (1990), which distinguishes between normative rationality and adaptive rationality (Oaksford and Chater, 1994, p. 608). Anderson treats the "General Principle of Rationality"—namely, the idea that "the cognitive system operates at all times to optimize the adaptation of the behavior of the organism"—as a scientific hypothesis for organizing the data regarding human cognition (Anderson, 1990, 28–9). Anderson (1990, p. 29) provides an iterative process for applying the General Principle of Rationality. The process requires researchers to consider the goals of the cognitive system, develop a formal model of the system's environment, and make minimal assumptions about the computational limitations of the cognitive system. With these features in place, researchers then determine the optimal behavior of the cognitive system in order to make testable, empirical predictions.

In their rational analysis of the WST, Oaksford and Chater (1994) treat maximizing information gain to be the goal of the cognitive systems involved in the selection task. Using Bayesian information theory (Shannon, 1948; Bar-Hillel and Carnap, 1953; Good, 1966), they then model the WST as a task of optimal data selection. Oaksford and Chater's rational analysis is a test of "Descriptive Bayesianism," the hypothesis that humans behave cognitively as if their aim was to obey Bayesian norms (Douven, 2022, p. 79). With respect to the WST, testing Descriptive Bayesianism requires determining whether respondents act as if they are Bayesian data selectors; put differently, Descriptive Bayesianism implies that researchers can use Bayesian epistemology to explicate reasoning.

Armed with information gain as the goal of respondents to the WST and Bayesian information theory as their formal framework, Oaksford and Chater specify minimal assumptions regarding the computational constraints on the cognitive system. This particular rational analysis relies on what Oaksford and Chater

[22]It should now be clear that the WST is closely connected to debates in the philosophy of science. The standard selection of cards espoused in the old, logical paradigm closely parallels Popper (1959a)'s falsificationist approach. Oaksford and Chater (2007, pp. 168-9) take this parallel to illustrate ways in which the philosophy of science influences the study of cognition.

(1994, p. 612) call the "rarity assumption." According to the rarity assumption, "Most objects, events and properties in the world are rare with respect to their complements—the things they are not" (Manktelow, 2012, p. 127). For the WST, this means that, according to Oaksford and Chater, respondents treated cards p showing **A** and q showing **2** as rare in comparison to their complements (i.e., $\neg p$ and $\neg q$, respectively). At first, the rarity assumption may seem implausible. How are respondents in WST judging the rarity of the cards they see?

Crucially, the claim of the rarity assumption is *not* that the respondents assume the **A** card (*qua* **A** card) and the **2** card (*qua* **2** card) to be rare; instead, the claim is that the cards playing the *roles* of the true antecedent p and true consequent q cards are assumed to be rare in comparison to the false antecedent $\neg p$ and false consequent $\neg q$ cards. Underpinning the assumption is the idea that the conditional rule tested in the WST communicates some sort of information about rare events. Manktelow (2012, p. 53) first presents the rarity assumption by saying that it "comes from the fact that most of the terms in natural language refer only to rare items and properties: the number of archbishops is a tiny proportion of the number of priests or believers, let alone men, humans, or living things." This presentation appears to tie the rarity assumption to the specificity of the terms used: more specific terms are rarer. Note, however, that this presentation of the assumption may not rest easily with the original conditional rule tested by Wason (i.e., *If a card has a vowel on one side, it has an even number on the other*). While vowels are indeed rarer than all letters, famous results of mathematics show that even integers are just as common as all integers (see Ewald (2023, p. 920–22)).

A better justification can be found in McKenzie et al. (2001, p. 308), which conducted studies indicating "that people have a strong tendency to spontaneously phrase conditional hypotheses in terms of rare, rather than common, events." For example, McKenzie et al. (2001) found that when asked to use conditionals to summarize sample data, participants were more likely to say, "If applicants have high SAT scores, then they will be accepted" rather than saying, "If applicants have average SAT scores, then they will be rejected." In this case, both high SAT scores and acceptance were rare, while average SAT scores and rejection were common. Because people do tend to use conditional hypotheses to communicate about rare events, it stands to reason that the respondents in the WST treated the conditional rule as information about the rarity of the p and q cards. Thus, the rarity assumption provided Oaksford and Chater with a putatively accurate model of the computational constraints involved in the WST.

Oaksford and Chater (1994) demonstrate that if the rarity assumption holds, then **A** (i.e., p) is the most informative card and **2** (i.e. q) is the second-most informative card. Of course, this ordering aligns with the selections observed in the task. The goal of maximizing information gain, the Bayesian model provided by information theory, and the rarity assumption coincide to reinstate the rationality of WST respondents. When the WST is conceived as a probabilistic inference

task, "the predominant response Wason and colleagues had found was in fact the one we should expect to find if people followed something close to an optimal Bayesian strategy for acquiring information" (Douven, 2022, p. 80). Here, Descriptive Bayesianism succeeds where Descriptive Deductivism failed.

As mentioned above, Oaksford and Chater's rational analysis of the WST is one small part of the new paradigm. (Oaksford and Chater, 2007) provides rational analyses of conditional and syllogistic reasoning as well. Moreover, Evans and Over (2004) and Douven (2015) provide Bayesian accounts of the conditional of their own. The account given by Evans and Over relies on the "Ramsey test", which uses the definition of **Conditional Probability** (see §1.1) as the basis for conditional reasoning. Since these works, the new paradigm has gained more traction in cognitive science; see (Elqayam and Over, 2013) for a special issue on the topic.

Two aspects of the new paradigm deserve attention prior to concluding our discussion. First, recognize that some consider probability—interpreted as degrees of belief or confidence—to be one of *three* pillars of the new paradigm (Bonnefon, 2013, p. 395). While the use of probability in rational analysis is a unifying theme across the new paradigm, utility theory and dual-process theory also play vital roles in the paradigm (Bonnefon, 2013; Elqayam and Over, 2013; Manktelow, 2012). Here, we limit our focus to the role played by probability due to the constraints and emphasis of this chapter, but that limited focus should not be taken as a dismissal of the roles played by utility theory or dual-process theory.[23]

Furthermore, the new paradigm uses Bayesian epistemology in a way that is markedly different from Tversky and Kahneman's HBP. While Tversky and Kahneman borrowed Bayesianism from epistemology and used it as a benchmark for testing failures in human reasoning, the new paradigm treats Bayesian epistemology as a bona fide hypothesis about the human mind. The new paradigm's reliance on Descriptive Bayesianism reconceives the relationship between probability theory and cognitive science, which in turn provides a new perspective on some classic findings from the HBP (e.g., see Griffiths and Tenenbaum (2006)). Prior to the new paradigm, Bayesian epistemology was treated in a purely normative way. Following the new paradigm, Bayesian epistemology operates as a bridge between cognitive science and philosophy.

2.2 Bayesian Statistics in Cosmology

As with its presence in the social sciences, so Bayesianism has come to be ubiquitous in today's hard sciences. In particular, in data analysis and statistics, Bayesian techniques are now used across the sciences, sometimes in strong preference to

[23]However, insofar as the recent advances in probabilistic, computational models of the mind (e.g., Griffiths et al. (2010, 2023)) are grouped under the "new paradigm" moniker, the use of Bayesian probability is arguably the only unifying theme across the new paradigm.

classical-Frequentist methods. We focus this section on the use of Bayesian statistics in modern cosmology, where such methods have a central role in investigating several distinct, crucial research questions.

Nearly two decades ago now, Trotta (2008, p. 71) already observed that "The application of Bayesian methods in cosmology and astrophysics has flourished over the past decade [and these methods] have proven to be vastly superior to more traditional statistical tools." The general framework for these various Bayesian applications may be summarized as follows. The cosmologist reasons through the lens of one or more of possible, alternative cosmological "scenarios" or models $\mathcal{M}_i$. Any such $\mathcal{M}_i$ is itself a collection of more specific hypotheses, these being sufficiently detailed to imply constrained values of cosmological parameters. A cosmological parameter typically represents "some physically meaningful quantity, such as the mass of an extra-solar planet or the abundance of dark matter in the Universe" (Trotta, 2008, p. 78). Formally then, such a hypothesis is a "vector of parameters." Modern cosmology's default model is the "vanilla" Lambda Cold Dark Matter model [abbreviated ΛCDM] (Trotta, 2008, §5.1). ΛCDM is also known as the "cosmological concordance model," since it posits a cosmological scenario describing the concordance of many commitments of the current scientific paradigm—an expanding universe, the initial singularity of the Big Bang, General Relativity, an isotropic and homogenous universe, etc.

Classical-Frequentist statistical methods would of course render the probability of ΛCDM or any other cosmological model $\mathcal{M}_i$ as nonsensical, insofar as it is nonsense to imagine the relative frequency of some $\mathcal{M}_i$ being instantiated in a repeatable sequence of universes. The Bayesian interpretation of probabilities, by contrast, counts $\Pr(\mathcal{M}_i)$ as a meaningful probability. Across different contexts of research, such prior probabilities might be assigned based on past, relevant studies and information; more commonly, in cases of model comparison, these are set equal across candidate models in line with **Indifference**'s injunction to assign equal priors when there is insufficient reason to do otherwise. Cosmologists occasionally defer to philosophical defenses of such assignments; e.g., Trotta (2008, p. 77) leans on "convergence theorems" for **Bayes's Rule** (see our discussion in §1.2 above) when he writes, "the prior assignment eventually becomes irrelevant as better and better data make the posterior likelihood-dominated."

New "cosmological observations" and data arise from various "probes," including the observations of distributions of galaxy clusters across time and measurements of "shears" through gravitational lensing (Trotta 2008, §5.2; Massimi 2021, p. 3). How well such varied data d fit with a particular model is summarized in the corresponding likelihood, referred to as the "Bayesian evidence" relative to $\mathcal{M}_i$ (Trotta, 2008, p. 81):

$$\Pr(d|\mathcal{M}) = \int_{\Omega_{\mathcal{M}}} \Pr(d|\theta \wedge \mathcal{M}) \Pr(\theta|\mathcal{M}) d\theta$$

Here, $\Pr(d|\mathcal{M})$ is assessed by averaging model $\mathcal{M}$'s various sub-hypotheses' like-

lihoods apropos d over the (in this case, continuous) parameter space $\Omega_{\mathcal{M}}$. Like $\Pr(\mathcal{M}_i)$, the Bayesian evidence relative to cosmological model $\mathcal{M}_i$ is plausibly only rendered meaningful on Bayesian grounds; this is at least due to the fact that values of $\Pr(\theta|\mathcal{M})$—corresponding to the prior probabilities of $\mathcal{M}_i$'s various sub-hypotheses—will generally be meaningless on the Frequentist interpretation.

Within this basic framework, cosmologists apply Bayesian methods to a variety of different questions and tasks. Following the helpful leads of both (Trotta, 2008) and (Massimi, 2021), we focus our discussion on three of these in particular: model comparison, cross-checking of data, and parameter inference. Bayesianism is applicable in interestingly distinct ways across these three areas of research.

The notion of Bayesian evidence has its most prominent, direct role in cosmological *model comparison*. Of course, within a Bayesian approach, the ultimate marker of which of competing models is better off in light of new observations and data d is a comparison of posterior probabilities, conditional on d. Formally, the ratio of posterior probabilities is a useful metric for assessing this comparison, as $\Pr(d)$ in the use of **Bayes's Theorem** cancels out in the calculation of this ratio:

$$\frac{\Pr(\mathcal{M}_1|d)}{\Pr(\mathcal{M}_2|d)} = \frac{\Pr(\mathcal{M}_1)\Pr(d|\mathcal{M}_1)\Pr(d)}{\Pr(\mathcal{M}_2)\Pr(d|\mathcal{M}_2)\Pr(d)} = \frac{\Pr(\mathcal{M}_1)}{\Pr(\mathcal{M}_2)} \cdot \frac{\Pr(d|\mathcal{M}_1)}{\Pr(d|\mathcal{M}_2)}.$$

In practice when it comes to model comparison, as mentioned above, it is standard for cosmologists to be "non-committal," by enforcing **Indifference** and thereby assuming equal prior probabilities across models. Consequently, such comparisons ultimately amount to direct comparisons of the Bayesian evidence relative to the models in question (Trotta 2008, p. 81; Massimi 2021, pp. 5-6):

$$\frac{\Pr(\mathcal{M}_1|d)}{\Pr(\mathcal{M}_2|d)} = \frac{\Pr(d|\mathcal{M}_1)}{\Pr(d|\mathcal{M}_2)} = \frac{\int_{\Omega_{\mathcal{M}_1}} \Pr(d|\theta \wedge \mathcal{M}_1)\Pr(\theta|\mathcal{M}_1)d\theta}{\int_{\Omega_{\mathcal{M}_2}} \Pr(d|\theta \wedge \mathcal{M}_2)\Pr(\theta|\mathcal{M}_2)d\theta}.$$

In this way, cosmological model comparison standardly comes down to a comparison of model likelihoods.

As data is gathered across diverse probes, Bayesian methods have a distinct, crucial role in the *cross-checking of data* for the sake of testing that data set's concordance or discordance on various models. Various formal, Bayesian criteria for such cross-checks are used by cosmologists in practice (Charnock et al. 2017, §II; Handley and Lemos 2019). To take one example, discussed at length by Handley and Lemos (2019) and by Massimi (2021), the "R statistic" effectively compares the joint probability assigned conditional on $\mathcal{M}_i$ to the conjunction of the different data—say d_1 and d_2—to what that joint probability would equal were d_1 and d_2 entirely independent of one another given $\mathcal{M}_i$:

$$\frac{\Pr(d_1 \wedge d_2|\mathcal{M}_i)}{\Pr(d_1|\mathcal{M}_i)\Pr(d_2|\mathcal{M}_i)}.$$

If $\mathcal{M}_i$ renders d_1 and d_2 incompatible, then the R statistic takes its minimum value 0. If $\mathcal{M}_i$ renders d_1 and d_2 statistically irrelevant, then the R statistic takes value 1. $0 < R < 1$ thus marks data that are discordant (though compatible) on M_i, while data are gauged as concordant to the extent that $R > 1$.

Beyond this, the interpretation and importance of the R statistic is a matter of some controversy amongst cosmologists. Handley and Lemos (2019, p. 3) remark that a common interpretation amongst cosmologists is to read the statistic as the "ratio of probabilities that the shared model parameters come from different universes in comparison with the probability that they come from the same universe." However, Handley and Lemos (2019, p. 3) go on to argue that this interpretation suggests a misleading, dubious relevance to the inferential task of model comparison (cf., Massimi 2021, pp. 7-8).

Applying **Conjunction** (§1.1) to the numerator of the R statistic, Handley and Lemos observe that it is equivalent to both of the following ratios:

$$\frac{\Pr(d_1 \wedge d_2|\mathcal{M}_i)}{\Pr(d_1|\mathcal{M}_i)\Pr(d_2|\mathcal{M}_i)} = \frac{\Pr(d_1|d_2 \wedge \mathcal{M}_i)}{\Pr(d_1|\mathcal{M}_i)} = \frac{\Pr(d_2|d_1 \wedge \mathcal{M}_i)}{\Pr(d_2|\mathcal{M}_i)}.$$

This formal result clarifies that R measures how much more expected one set of data becomes in virtue of already accepting the other and through the lens of model $\mathcal{M}_i$ (Handley and Lemos, 2019, p. 3). Informally, R measures the degree to which the distinct data hang together in light of the model. Here, it's worth briefly pointing out another direct point of connection between formal epistemology and this scientific application of Bayesianism. In fact, the R statistic corresponds exactly to a measure of the epistemological concept of *coherence* offered in Bayesian epistemology, most famously by Shogenji (1999). Formal epistemological research on this measure thus bears directly on the critical evaluation of R—e.g., Schupbach (2011) provides an argument for how properly to generalize Shogenji's measure of coherence [R] to the case where we are assessing $n > 2$ data sets.

Bayesian methods also have direct application in cosmology when it comes to *parameter inference* (Trotta, 2008, §3). In this context, the cosmologist assumes a cosmological model—e.g., the default ΛCDM—and then aims to use dataset d to tune and further constrain the parameters of that model. In choosing between different possible values for a particular parameter θ, the cosmologist once again makes use of posterior probabilities. **Bayes's Theorem** has one calculate the probability of this parameter taking a particular value θ_i, conditional on d and model $\mathcal{M}$, as follows:

$$\Pr(\theta_i|d \wedge \mathcal{M}) = \frac{\Pr(\theta_i|\mathcal{M})\Pr(d|\theta_i \wedge \mathcal{M})}{\Pr(d|\mathcal{M})}.$$

The denominator term $\Pr(d|\mathcal{M})$ is the Bayesian evidence, which is constant across any such calculation for parameter settings. As such, it has no distinguishing effect on the various possible values for θ and so may be ignored with the requisite

calculation more simply being:

$$\Pr(\theta_i|d \wedge \mathcal{M}) \propto \Pr(\theta_i|\mathcal{M}) \Pr(d|\theta_i \wedge \mathcal{M}).$$

Bayesian parameter inference, model comparison, and data cross-checking all involve idiosyncratically Bayesian probabilities. None of these tasks, as described above, may be pursued, or even made sense of, from a Classical-Frequentist perspective. And each task is distinct, revealing different ways in which Bayesianism is applicable.

Massimi (2021) argues that Bayesianism is applicable in *problematically* distinct ways across these tasks, insofar as the very same prior probabilities are assigned differing values, encapsulating varying amounts of informativeness. As Massimi (2021, p. 8) writes, "Statistically, one cannot use different priors for different problems concerning the same data and the same parameters, especially since the priors at stake here are theory-dependent and data-dependent, but they are not subjective degrees of belief of cosmologist A vs cosmologist B." We close this section with an all-too-brief discussion of this criticism, as it allows us to highlight just how far removed from its philosophical foundations the modern practice of Bayesian statistics can become.

As we have seen, in the traditional view, probability theory is interpreted as providing an inductive logic of *degrees of belief*. That said, when used as a feature of statistical analysis, general to a field and research question, the use and interpretation of these probabilities apparently becomes more fluid. Many modern self-styled Bayesian statisticians want little to do with Bayesianism's semantic and philosophical foundations. As Gelman and Shalizi (2013) describe it, "the practice of Bayesian statistics" is better viewed as a model-building and checking endeavor, having much in common with the "hypothetico-deductive" practice of Classical-Frequentist statistics (p. 10).

Insofar as Gelman and Shalizi are correct in their description of the common practice of Bayesian statistics today, most Bayesian statisticians are aptly categorized as "instrumentalist Bayesians." Unlike traditional Bayesians, who offer a conceptual and philosophical grounding for their liberal use of probabilities (including prior probabilities of hypotheses), the instrumentalist settles for a purely pragmatic, proof-in-the-pudding approach. Probability theory is not helpfully interpreted as an inductive logic for uncertainties, and the idea that probabilities are normative counterparts for graded, doxastic attitudes only constrains us from being able to wield them freely. Instead, probabilities and probabilistic models are instruments or tools for fitting data, predicting data, and so on (Gelman and Shalizi, 2013, p. 19).

Instrumentalist Bayesians willingly make use of full probability models as powerful instruments for fitting and predicting data but without necessarily accepting (or seeing the need to defend) traditional Bayesian foundations for such probabilistic application. What sets Bayesian statistics apart from Classical statis-

tics then, on this instrumentalist view, is *merely* the fact that Bayesian statistical models make use of "full joint distributions, embracing observables, latent variables, and parameters." To summarize the point curtly, instrumentalist Bayesians are statisticians who use Classical-Frequentist methods to do standard model-building and checking while helping themselves to models with more free parameters. They are Bayesians only in the sense that they dip into the Bayesian toolbox. But they use Bayesian models *sans* the Bayesian foundations, purely for pragmatic purposes.[24]

Let us return now to Massimi's concern that the same priors are evaluated differently across different scientific tasks in cosmology. This concern—coming from a philosopher as it were—may truly be empty to the pragmatist, instrumentalist Bayesian. Across different research questions, different formal tools may be variably useful. Insofar as these formal tools (e.g., prior probabilities of models) are not tied to any actual interpretation, there seems to be no motivation for using them consistently across different applications. Consistency be damned; better to use the tools in whatever way "works" best, peculiar to each task. While we sympathize with Massimi's worries, we also suggest that they are thus not so worrying from within the statistical practice of cosmology; at least this is the case accepting Gelman and Shalizi's insiders' perspective on the state of their own discipline. Massimi's concerns are rather more telling of how far from its philosophical-semantic foundations modern Bayesian statistical practice can become.

Conclusion

In the foreword to his seminal work at the foundations of Bayesianism, Ramsey (1926) is careful to distinguish his notion of probability "taken as a branch of logic, the logic of partial belief and inconclusive argument" from the notion of probability used "in statistical and physical science." As he writes, "Indeed the general difference of opinion between statisticians [and physicists] who for the most part adopt the frequency theory of probability and logicians who mostly reject it renders it likely that the two schools are really discussing different things."

Modern Bayesians, writing a century later, still treat probability as a logic of partial belief. Moreover, it has only become more clear in the last century that the theory of probability is of fundamental importance in the sciences. However, Ramsey was not altogether prescient in suggesting that Bayesianism would have

[24]Note that instrumentalist Bayesians thus fall outside the ranks of Bayesianism as we characterize it in §1.1, since we characterize Bayesians as being unified by their interpretation of probabilities. Are instrumentalist "Bayesians" truly worthy of the name then? We suspect that one gives away too much by allowing statisticians to categorize themselves as "Bayesian" simply by virtue of being willing to countenance models with certain extra parameters. Obviously, exploring and defending this claim would go far beyond the scope of this chapter.

its application in logic with Frequentism capturing the scientifically useful notion of probability. This chapter has illustrated ways in which scientists (from cognitive scientists to physicists) *integrate* the Bayesian logic of partial belief into their investigations of the world.

Thanks in large part to the rise of modern computing, Bayesianism has experienced a century of growth in the sciences. While that century has at times been tumultuous (e.g., the so-called "Statistics Wars" between Bayesians and Classical-Frequentists), the diverse applications of Bayesian methodology have indisputably generated illuminating discussion. Even for cases as diverse as adjudicating the rationality of judgments (§2.1) and comparing cosmological models of the universe (§2.2), Bayesianism has provided scientists with a common currency for reasoning under uncertainty. There is no telling what the next century holds for Bayesianism in science; we have high degrees of belief that it will be fascinating to watch unfold.

Further Reading

Bayesianism Across the Sciences:

- Biology: Mann et al. (2022)
- Econometrics: West and Harrison (1997); Geweke (2005); Martin et al. (2024)
- Evolutionary Game Theory: Gintis (2009)
- Precision Agriculture: Sivaraj and Palanisamy (2022); Ribeiro et al. (2023)

Popular Press Treatments of Bayesianism: McGrayne (2011); Pearl and Mackenzie (2020); Silver (2020)

Philosophical Introductions to Bayesianism: Resnik (1987); Kadane et al. (1999); Sprenger and Hartmann (2019)

Philosophy of Probability: Schum (1994); Leitgeb (2014); Bandyopadhyay et al. (2016); Peden (2021); Hájek (2023)

Alternatives to Bayesianism: Royall (1997); Mayo (1996, 2018); Glymour (2001); Fletcher and Mayo-Wilson (2023)

References

Anderson, J. R. (1990). *The Adaptive Character of Thought*. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Bandyopadhyay, P. S., Brittan, G., and Taper, M. L. (2016). *Belief, Evidence, and Uncertainty: Problems of Epistemic Inference*. SpringerBriefs in Philosophy. Springer International Publishing, Cham.

Bar-Hillel, Y. and Carnap, R. (1953). Semantic information. *The British Journal for the Philosophy of Science*, IV(14):147–157.

Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society*, 53:370–418.

Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Oxford University Press, Incorporated, Oxford, United Kingdom.

Bertrand, J. (1889). *Calcul des Probabilités*. Gauthier-Villars, Paris.

Blackwell, D. and Dubins, L. (1962). Merging of opinions with increasing information. *The Annals of Mathematical Statistics*, 33(3):882–886.

Bonnefon, J.-F. (2013). New ambitions for a new paradigm: Putting the psychology of reasoning at the service of humanity. *Thinking & Reasoning*, 19(3-4):381–398.

Bottemanne, H. (2025). Bayesian brain theory: Computational neuroscience of belief. *Neuroscience*, 566:198–204.

Bovens, L. and Hartmann, S. (2002). Bayesian networks and the problem of unreliable instruments. *Philosophy of Science*, 69(1):29–72.

Bovens, L. and Hartmann, S. (2003). *Bayesian Epistemology*. Oxford University Press, New York.

Brekel, J. (2025). A new look at Keynesian weight of evidence. *Erkenntnis*.

Carnap, R. (1962). *Logical Foundations of Probability*. University of Chicago Press, Chicago, 2nd edition.

Charnock, T., Battye, R. A., and Moss, A. (2017). Planck data versus large scale structure: Methods to quantify discordance. *Physical Review D*, 95(123535):1–19.

Cohen, L. J. (1979). On the psychology of prediction: Whose is the fallacy? *Cognition*, 7(4):385–407.

Cohen, L. J. (1981). Can Human Irrationality Be Experimentally Demonstrated? *Behavioral and Brain Sciences*, 4(3):317–331.

Cox, R. T. (1961). *The Algebra of Probable Inference*. The John Hopkins University Press, Baltimore.

Crupi, V., Elia, F., Aprà, F., and Tentori, K. (2018). Double conjunction fallacies in physicians' probability judgment. *Medical Decision Making*, 38(6):756–760. PMID: 29978726.

Crupi, V., Fitelson, B., and Tentori, K. (2008). Probability, confirmation, and the conjunction fallacy. *Thinking and Reasoning*, 14(2):182–199.

Crupi, V., Tentori, K., and Gonzalez, M. (2007). On Bayesian measures of evidential support: Theoretical and empirical issues. *Philosophy of Science*, 74:229–252.

de Finetti, B. (1937). La Prévision: ses Lois Logiques, ses Sources Subjectives. *Annales de l'institut Henri Poincaré*, 7:1–68. Translated by H. Kyburg as "Foresight: its Logical Laws, its Subjective Sources," in (Kyburg and Smokler, 1964, pp. 93-158).

Douven, I. (2015). *The Epistemology of Indicative Conditionals: Formal and Empirical Approaches*. Cambridge University Press, Cambridge.

Douven, I. (2022). *The Art of Abduction*. Massachusetts Institute of Technology Press, Cambridge, Massachusetts.

Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. MIT Press, Cambridge, Mass.

Easwaran, K. (2011). Bayesianism I: Introduction and arguments in favor. *Philosophy Compass*, 6(5):312–320.

Easwaran, K. (2019). Conditional probabilities. In Pettigrew, R. and Weisberg, J., editors, *The Open Handbook of Formal Epistemology*, pages 131–198. PhilPapers Foundation.

Edwards, W., Lindman, H., and Savage, L. J. (1963). Bayesian statistical inference for psychological research. *Psychological Review*, 70(3):193–242.

Elqayam, S. and Over, D. E. (2013). New Paradigm Psychology of Reasoning: An Introduction to the Special Issue Edited by Elqayam, Bonnefon, and Over. *Thinking & Reasoning*, 19(3-4):249–265.

Evans, J. S. B. T. and Over, D. E. (2004). *If*. Oxford Cognitive Science Series. Oxford University Press, Oxford ; New York.

Ewald, W. B., editor (2023). *From Kant to Hilbert: A Source Book in the Foundations of Mathematics; Volume II*. Oxford Scholarship Online. Clarendon, Oxford.

Fiedler, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research*, 50:123–129.

Fitelson, B. (1999). The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science*, 66:S362–S378.

Fitelson, B. (2001). A Bayesian Account of Independent Evidence with Applications. *Philosophy of Science*, 68(3,):S123–S140.

Fletcher, S. C. and Mayo-Wilson, C. (2023). *Evidence in Classical Statistics*, pages 515–527. Routledge, London, 1 edition.

Gaifman, H. and Snir, M. (1982). Probabilities over rich languages, testing and randomness. *The Journal of Symbolic Logic*, 47(3):495–548.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2014). *Bayesian Data Analysis*. CRC Press, Boca Raton, FL, 3rd edition.

Gelman, A. and Shalizi, C. R. (2013). Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, 66:8–38.

Geweke, J. (2005). *Contemporary Bayesian Econometrics and Statistics*. John Wiley & Sons.

Gigerenzer, G. (1991). How to Make Cognitive Illusions Disappear: Beyond "Heuristics and Biases". *European Review of Social Psychology*, 2(1):83–115.

Gigerenzer, G. and Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102(4):684–704.

Gilovich, T., Griffin, D. W., and Kahneman, D. (2002). *Heuristics and Biases: The Psychology of Intuitive Judgement*. University Press, Cambridge, 1st edition.

Gintis, H. (2009). *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. Princeton University Press, Princeton, 2nd ed edition.

Glass, D. H. (2023). How good is an explanation? *Synthese*, 201(53):1–26.

Glymour, C. (2001). Instrumental Probability. *The Monist*, 84(2):284–300.

Good, I. J. (1966). A derivation of the probabilistic explication of information. *Journal of the Royal Statistical Society: Series B (Methodological)*, 28(3):578–581.

Good, I. J. (1968). Corroboration, explanation, evolving probability, simplicity and a sharpened razor. *British Journal for the Philosophy of Science*, 19(2):123–143.

Good, I. J. (1971). 46656 varieties of Bayesians. *American Statistician*, 25:62–63.

Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., and Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences*, 14(8):357–364.

Griffiths, T. L., Kemp, C., and Tenenbaum, J. B. (2023). Bayesian Models of Cognition. In Sun, R., editor, *The Cambridge Handbook of Computational Cognitive Sciences*, pages 80–138. Cambridge University Press, 2nd edition.

Griffiths, T. L. and Tenenbaum, J. B. (2006). Optimal Predictions in Everyday Cognition. *Psychological Science*, 17(9):767–773.

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

Hájek, A. (2023). Interpretations of probability. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Winter 2023 edition.

Handley, W. and Lemos, P. (2019). Quantifying tensions in cosmological parameters: Interpreting the DES evidence ratio. *Physical Review D*, 100(043504):1–15.

Hintikka, J. (2004). A Fallacious Fallacy? *Synthese*, 140(1-2):25–35.

Hohwy, J. (2014). *The Predictive Mind*. Oxford University Press, Oxford.

Howson, C. and Urbach, P. (2006). *Scientific Reasoning: The Bayesian Approach*. Open Court, Peru, IL, 3rd edition.

Humphreys, P. (1985). Why propensities cannot be probabilities. *Philosophical Review*, 94(4):557–570.

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

Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106(4):620–630.

Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge.

Jeffrey, R. C. (1983). *The Logic of Decision*. University of Chicago Press, Chicago, 2nd edition.

Jeffreys, H. (1939). *Theory of Probability*. Clarendon Press, Oxford.

Jolls, C., Sunstein, C. R., and Thaler, R. (1998). A Behavioral Approach to Law and Economics. *Stanford Law Review*, 50(5):1471–1550.

Jönsson, M. L. and Shogenji, T. (2019). A unified account of the conjunction fallacy by coherence. *Synthese*, 196:221–237.

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

Kadane, J. B., Schervish, M. J., and Seidenfeld, T. (1999). *Rethinking the Foundations of Statistics*. Cambridge Studies in Probability, Induction and Decision Theory. Cambridge University Press, Cambridge.

Kahneman, D. (2011). *Thinking, Fast and Slow*. MRSV. Farrar, Straus and Giroux, New York, first edition.

Kahneman, D., Sibony, O., and Sunstein, C. R. (2022). *Noise: A Flaw in Human Judgment*. Little, Brown Spark, New York, NY Boston London, first edition.

Kahneman, D., Slovic, P., and Tversky, A., editors (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge ; New York.

Kahneman, D. and Tversky, A. (1996). On the Reality of Cognitive Illusions. *Psychological Review*, 103(3):582–591.

Kasser, J. (2016). Two Conceptions of Weight of Evidence in Peirce's Illustrations of the Logic of Science. *Erkenntnis*, 81(3):629–648.

Kemeny, J. G. and Oppenheim, P. (1952). Degree of factual support. *Philosophy of Science*, 19:307–324.

Keynes, J. M. (1921). *A Treatise on Probability*. Macmillan, London.

Kolmogorov, A. N. (1933). *Foundations of the Theory of Probability*. Chelsea Publishing Company, New York.

Kuhn, T. S. (2012). *The Structure of Scientific Revolutions*. The University of Chicago Press, Chicago ; London, fourth edition.

Kyburg, H. E. and Smokler, H. E., editors (1964). *Studies in Subjective Probability*. John Wiley, New York.

Laplace, P.-S. (1814). *A Philosophical Essay on Probabilities*.

Leitgeb, H. (2014). The Stability Theory of Belief. *The Philosophical Review*, 123(2):131–171.

Levi, I. (2004). Jaakko Hintikka. *Synthese*, 140:37–41.

Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85(3):297–315.

Lewis, M. (2017). *The Undoing Project: A Friendship That Changed Our Minds*. W.W. Norton & Company, New York London, first edition.

Lin, H. (2024). Bayesian epistemology. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2024 edition.

Luce, R. D. and Raiffa, H. (1957). *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons, New York.

Lyon, A. (2016). Kolmogorov's axioms and its discontents. In Hájek, A. and Hitchcock, C., editors, *The Oxford Handbook of Probability and Philosophy*, pages 155–166. Oxford University Press, Oxford.

Mangalam, M. (2025). The myth of the Bayesian brain. *European Journal of Applied Physiology*, 125(10):2643–2677.

Manktelow, K. (2012). *Thinking and Reasoning: An Introduction to the Psychology of Reason, Judgment and Decision Making*. 360 Degree Business. Psychology Press, Hove, East Sussex, 2nd edition.

Mann, S. F., Pain, R., and Kirchhoff, M. D. (2022). Free energy: A user's guide. *Biology & Philosophy*, 37(4):33.

Marr, D. (2010). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. The MIT Press.

Martin, G. M., Frazier, D. T., Maneesoonthorn, W., Loaiza-Maya, R., Huber, F., Koop, G., Maheu, J., Nibbering, D., and Panagiotelis, A. (2024). Bayesian forecasting in economics and finance: A modern review. *International Journal of Forecasting*, 40(2):811–839.

Massimi, M. (2021). Cosmic Bayes: Datasets and priors in the hunt for dark energy. *European Journal for Philosophy of Science*, 11(29):1–21.

Mayo, D. G. (1996). *Error and the Growth of Experimental Knowledge*. University of Chicago Press, Chicago.

Mayo, D. G. (2018). *Statistical Inference as Severe Testing: How to Get beyond the Statistics Wars*. Cambridge University Press, Cambridge; New York, NY.

McGrayne, S. B. (2011). *The Theory That Would Not Die: How Bayes' Rule Cracked the Enigma Code, Hunted down Russian Submarines, & Emerged Triumphant from Two Centuries of Controversy*. Yale university press, New Haven.

McKenzie, C. R., Ferreira, V. S., Mikkelsen, L. A., McDermott, K. J., and Skrable, R. P. (2001). Do Conditional Hypotheses Target Rare Events? *Organizational Behavior and Human Decision Processes*, 85(2):291–309.

Nisbett, R. E. (1980). *Human Inference: Strategies and Shortcomings of Social Judgment*. Century Psychology Series. Prentice-Hall, Englewood Cliffs, N.J.

Oaksford, M. and Chater, N. (1994). A Rational Analysis of the Selection Task as Optimal Data Selection. *Psychological Review*, 101(4):608–631.

Oaksford, M. and Chater, N. (2007). *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford University Press, Oxford.

Oaksford, M. and Chater, N. (2009). Précis of *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. *Behavioral and Brain Sciences*, 32(1):69–84.

Pearl, J. and Mackenzie, D. (2020). *The Book of Why: The New Science of Cause and Effect*. Basic Books, New York, first trade paperback edition edition.

Peden, W. (2021). Probability and arguments: Keynes's legacy. *Cambridge Journal of Economics*, 45(5):933–950.

Peirce, C. S. (1992). *The Essential Peirce: Selected Philosophical Writings*, volume 1. Indiana University Press, Bloomington.

Popper, K. R. (1959a). *The Logic of Scientific Discovery*. Hutchinson, London.

Popper, K. R. (1959b). The propensity interpretation of probability. *British Journal for the Philosophy of Science*, 10(37):25–42.

Putnam, H. (1981). *Reason, Truth and History*. Cambridge University Press, New York.

Ramsey, F. P. (1926). Truth and Probability. In Mellor, D. H., editor, *Philosophical Papers*, pages 52–94. Cambridge University Press, Cambridge. Edited collection published in 1990.

Redelmeier, D. A., Rozin, P., and Kahneman, D. (1993). Understanding Patients' Decisions: Cognitive and Emotional Perspectives. *JAMA*, 270(1):72–76.

Resnik, M. D. (1987). *Choices: An Introduction to Decision Theory*. University of Minnesota Press, Minneapolis, MN.

Ribeiro, V. P., Desuó Neto, L., Marques, P. A. A., Achcar, J. A., Junqueira, A. M., Chinatto, A. W., Junqueira, C. C. M., Maciel, C. D., and Balestieri, J. A. P. (2023). A Stochastic Bayesian Artificial Intelligence Framework to Assess Climatological Water Balance under Missing Variables for Evapotranspiration Estimates. *Agronomy*, 13(12):2970.

Rosenkrantz, R. D. (1981). *Foundations and Applications of Inductive Probability*. Ridgeview Press, Atascadero, CA.

Royall, R. M. (1997). *Statistical Evidence: A Likelihood Paradigm*. Number 71 in Monographs on Statistics and Applied Probability. Chapman & Hall, Boca Raton, repr edition.

Savage, L. J. (1954). *The Foundations of Statistics*. Wiley, New York. Pages references are to the 1972 Dover edition.

Schippers, M. and Siebel, M. (2021). The logic of coherence. *Synthese*, 198(8):7697–7714.

Schum, D. A. (1994). *The Evidential Foundations of Probabilistic Reasoning*. Wiley Series in Systems Engineering.

Schupbach, J. N. (2011). New hope for Shogenji's coherence measure. *British Journal for the Philosophy of Science*, 62:125–142.

Schupbach, J. N. (2012). Is the conjunction fallacy tied to probabilistic confirmation? *Synthese*, 184(1):13–27.

Schupbach, J. N. (2022). *Bayesianism and Scientific Reasoning*. Elements in the Philosophy of Science. Cambridge University Press, Cambridge.

Schupbach, J. N. (2023). On the logical structure of best explanations. *Philosophy of Science*, 90(5):1150–1160.

Schupbach, J. N. and Glass, D. H., editors (2023). *Conjunctive Explanations: The Nature, Epistemology, and Psychology of Explanatory Multiplicity*. Routledge Studies in the Philosophy of Science. Routledge, New York.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656.

Shogenji, T. (1999). Is coherence truth-conducive? *Analysis*, 59(4):338–345.

Silver, N. (2020). *The Signal and the Noise: Why so Many Predictions Fail - but Some Don't*. A Penguin Book Economics/Politics/Sports. Penguin Books, New York, NY, published with a new preface in penguin books 2020 edition.

Sivaraj, A. and Palanisamy, V. (2022). Deep neural network based interactive fuzzy Bayesian search algorithm for low-cost smart farming automation model. *Concurrency and Computation: Practice and Experience*, 34(27):e7296.

Sprenger, J. and Hartmann, S. (2019). *Bayesian Philosophy of Science*. Oxford University Press, Oxford.

Stanovich, K. E. (1999). *Who Is Rational?: Studies of Individual Differences in Reasoning*. Lawrence Erlbaum Associates, Mahwah, N.J.

Tentori, K. and Crupi, V. (2012). How the conjunction fallacy is tied to probabilistic confirmation: Some remarks on Schupbach (2012). *Synthese*, 184(1):3–12.

Tentori, K., Crupi, V., and Russo, S. (2013). On the determinants of the conjunction fallacy: Probability versus inductive confirmation. *Journal of Experimental Psychology: General*, 142(1):235–255.

Thaler, R. (1980). Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior & Organization*, 1(1):39–60.

Thaler, R. H. and Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press, New Haven, USA, new international edition.

Titelbaum, M. G. (2013). *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford University Press, Oxford.

Titelbaum, M. G. (2022). *Fundamentals of Bayesian Epistemology 2: Arguments, Challenges, Alternatives*. Oxford University Press, Oxford.

Trotta, R. (2008). Bayes in the sky: Bayesian inference and model selection in cosmology. *Contemporary Physics*, 49(2):71–104.

Tversky, A. and Kahneman, D. (1974). Judgment Under Uncertainty: Heuristics and Biases. *Science*, 185(4157):1124–1131.

Tversky, A. and Kahneman, D. (1983). Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment. *Psychological Review*, 90(4):293–315.

van Fraassen, B. C. (1989). *Laws and Symmetry*. Oxford University Press, New York.

von Mises, R. (1957). *Probability, Statistics, and Truth*. Dover, Mineola, NY.

von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.

Wason, P. C. (1968). Reasoning About a Rule. *Quarterly Journal of Experimental Psychology*, 20(3):273–281.

Wasserman, L. (2004). *All of Statistics: A Concise Course in Statistical Inference*. Springer, New York.

Weisberg, J. (2009). Varieties of Bayesianism. In Gabbay, D. M., Hartmann, S., and Woods, J., editors, *Handbook of the History of Logic*, volume 10: Inductive Logic, pages 477–551. North-Holland, Amsterdam.

West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models*. Springer.

Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford University Press, Oxford.