# Bayesian Cognitive Science, Monopoly, and Neglected Frameworks

December 15, 2016

**Abstract**

Cognitive scientists, and a number of philosophers also, widely agree that discovering and assessing explanations of cognitive phenomena involving states of uncertainty should be done in a Bayesian framework. One assumption in support of the modelling choice is that Bayes provides the *best* method for representing uncertainty. It is unclear, though, that Bayes possesses special epistemic virtues over alternative modelling frameworks, especially in explaining cognitive states of uncertainty since a systematic comparison has yet to be attempted. Currently, it is premature to assert that cognitive phenomena involving uncertainty are *best* explained within the Bayesian framework. The forewarning given in this paper is that progress in understanding how minds grapple with uncertainty may be hindered if too many philosophers and scientists continue to focus their efforts on Bayesian modelling, which risks monopolizing resources that may be better allocated to alternative approaches.

**Keywords**: Bayesian cognitive science, representing uncertainty, scientific realism, underdetermination thesis

## 1 Introduction

Bayesianism has become ever more prominent the cognitive and brain sciences as well as philosophy.[1] Driven by mathematical advances in statistics and computer

[1] The term 'Bayesianism' will be used as a placeholder for a diverse set of interrelated principles, methods, and problem-solving procedures, which are unified by three core ideas. First: single real-valued probability distributions ought to be used to represent uncertainty. Second: our degrees of belief at a given time ought to satisfy the axioms of the probability calculus. Third: degrees of belief, represented by determinate probabilities, ought to be updated "optimally" in

science, as well as engineering successes in fields such as machine learning and artificial intelligence, Bayesian modelling has been used in studying many cognitive phenomena including perception, motor control, learning, decision-making, and reasoning (Chater, Tenenbaum and Yuille 2006; Doya et al. 2007; Gopnik et al. 2004; Knill and Richards 1996; Körding 2007; Maloney 2002; Rao, Lewicki, and Olhausen 2002; Oaksford and Chater 2007; Tenenbaum et al. 2011).

A common reason for favoring Bayesianism is given by the fact that uncertainty is an ineliminable feature of cognitive systems' interactions with the world. In order to survive and behave adaptively, biological cognitive systems must rely on knowledge derived from sparse, noisy, and ambiguous sensory data produced by a constantly changing environment. Because sensory data are sparse, ambiguous, and corrupted by noise, cognitive systems constantly face problems of inference and decision making under uncertainty. Unless these problems are effectively solved, reliable action, accurate perception, and adaptive learning would not be achievable.

Provided that uncertainty is an ineliminable feature of cognitive systems' interactions with the world—so continues the argument—the explanatory framework cognitive scientists and philosophers use to seek explanations of epistemic and cognitive phenomena should account for how cognitive systems effectively deal with uncertainty. The framework ought to include sound inferential and decision-making procedures in light of uncertainty, which biological cognitive systems might deploy when interacting with their environment. Because the Bayesian framework is the best (allegedly) for representing uncertainty and providing such procedures, it ought to be the default framework used in seeking explanations of how cognitive systems handle uncertainty by making sound inferences and decisions in their local environments.

The argument just canvassed will be referred to as the *argument from uncertainty for Bayesian cognitive science*. The first aim of this paper is to carefully reconstruct it. This reconstruction will contribute to clarify the conceptual foundations of Bayesianism in relation to cognitive modelling. While philosophers have devoted a great deal of attention to Bayesianism as a model of rational belief and rational choice, they have made little effort to elucidate the nature, aims, scope, and implications of Bayesian modelling in cognitive science. This is unfortunate given the prominence of Bayesianism in cognitive science, and given recent controversy among cognitive scientists about both the empirical and theoretical adequacy of Bayesianism (e.g., Bowers and Davis 2012a, 2012b; Chater et al. 2011; Grif-

light of new data. The most familiar updating rule within Bayesian epistemology and cognitive science is simple Bayesian conditionalisation. However, Bayesians rely also on other rules for learning and inference, including Jeffrey's rule of conditionalisation (Zhao and Osherson 2010), the minimization of the free energy of sampled data (Friston 2009), the minimization of the Kullback-Leibler divergence (Diaconis and Zabell 1982; Hartmann and Rafiee Rad 2016), and Monte Carlo methods (Sanborn, Griffiths, and Navarro, 2010) fiths et al. 2012; Jones and Love 2011; for recent philosophical discussions see Colombo and Seriès 2012; Colombo and Hartmann 2015; Eberhardt and Danks 2011; Rescorla 2015a).

The second aim of our paper is to explicate under which conditions many cognitive scientists' decision to work within a Bayesian framework is justified. We believe that ultimately this issue will be resolved empirically, and in fact many debates about Bayesianism in cognitive science revolve around the question of whether some non-Bayesian model is equally good at fitting data. However, an important, neglected aspect of cognitive scientists' decision to work within Bayesian concerns the non-empirical (or superempirical) virtues of Bayesianism. Our contribution fills this gap by examining whether the argument from uncertainty provides sufficient justification for choosing Bayes as a modelling framework in cognitive science.

In doing so, we first clarify what the argument is supposed to establish (Section 2) and what role it plays in current scientific practice (Section 3). Then, we consider whether there is any genuine alternative to Bayesianism as a framework for representing uncertainty and for explaining how biological cognitive systems effectively deal with it. If alternatives are available, but happen to be systematically neglected, then the strength of the argument from uncertainty is weakened (Section 4). The upshot is that Bayes does not enjoy special non-empirical virtues in comparison to alternatives.

In the conclusion of the paper, a practical reason for choosing a Bayesian approach is given. Specifically, the Bayesian framework in comparison to alternatives currently affords cognitive scientists with a richer body of knowledge and tools that have been well-developed and employed in neighboring fields of machine learning, artificial intelligence, and statistics (Jordan and Mitchell 2015; Ghahramani 2015). Provided this reason, the popularity of Bayesianism indicates that theory choice in cognitive science, similarly to other scientific fields, is not often guided by truth, or by super-empirical virtues like simplicity, elegance, and unifying power. Instead, cognitive scientists choose a theory that is most compatible with available tools and methods from neigbouring fields (Gigerenzer 1991). However, this practice triggers the concern that progress in cognitive science may be hindered if too many cognitive scientists continue to invest their efforts solely in a single approach, which ultimately risks monopolizing resources that may be better allocated to alternative approaches.

## 2 From Uncertainty to Bayesian Brains

The *argument from uncertainty* seeks to provide a compelling reason for choosing the Bayesian framework in the attempt to explain cognitive phenomena whose production requires cognitive systems to handle uncertainty. The argument involves two steps. The first step aims to substantiate the claim that biological cognitive systems must effectively deal with uncertainty in order to survive and thrive (i.e. in order to interact adaptively with their environment). The second step aims to establish that Bayesianism is the best for explaining how cognitive systems effectively deal with uncertainty.

## 2.1 Uncertainty, Underdetermination, and Noise

Most Bayesian cognitive scientists introduce their studies by pointing out that the mind must constantly grapple with uncertainty. For example, Knill and Pouget (2004) introduce their discussion of the Bayesian brain hypothesis by claiming that "humans and other animals operate in a world of sensory uncertainty" (Knill and Pouget 2004, p. 712). Ma et al. (2006) motivate their study on how populations of neurons might perform Bayesian computations by saying that "virtually all computations performed by the nervous system are subject to uncertainty" (Ma et al. 2006, p. 1432). Oaksford and Chater (2007) advocate a Bayesian approach to human rationality suggesting that "human reasoning is well-adapted to the uncertain character of everyday reasoning to integrating and applying vast amounts of world knowledge concerning a partially known and fast-changing environment" (Oaksford and Chater 2007, p. 67). Pouget et al. (2013, p. 1170) write that "uncertainty is an intrinsic part of neural computation, whether for sensory processing, motor control or cognitive reasoning." Orbán and Wolpert (2011) focus on Bayesian approaches to sensorimotor control and motivate their focus by saying that "uncertainty is ubiquitous in our sensorimotor interactions, arising from factors such as sensory and motor noise and ambiguity about the environment" (Orbán and Wolpert 2011, p. 1). Tenenbaum et al. (2011) point out that "we build rich causal models, make strong generalizations, and construct powerful abstractions, whereas the input data are sparse, noisy, and ambiguous—in every way far too limited" (Tenenbaum et al. 2011, p. 1279). Vilares and Körding (2011) review several lines of research in Bayesian cognitive neuroscience emphasizing that "uncertainty is relevant in most situations in which humans need to make decisions and will thus affect the problems to be solved by the brain" (Vilares and Körding 2011, p. 22).

There are two ways of understanding the emphasis on uncertainty when it comes to motivating the Bayesian approach in cognitive science. According to one reading, the point is merely expository. Claims like the ones just quoted get readers to see why the Bayesian approach is worth considering when one is interested in understanding how cognitive systems deal with uncertainty. According to another reading, the point is justificatory. These claims serve to convince readers that cognitive scientists' choice to pursue the Bayesian approach is justified.

Granted, cognitive systems must grapple with uncertainty. But an inference from 'biological agents are faced with states of uncertainty' to 'Bayes is the best' is invalid on either point. If the point is merely expository, then it is puzzling that no alternatives to the Bayesian approach are ever mentioned in the literature. Bayesianism is just one possible approach for representing and handling inference under uncertainty, but it is by no means the only one (Halpern 2003). While cognitive scientists could choose some other non-Bayesian approach, the alternative theories just so happen to get neglected without apparent reason.

On the other hand, if the point is justificatory, then there is an implicit assumption that Bayesianism is somehow more empirically adequate, or enjoys greater (or more) non-empirical virtues over alternative frameworks. However, such an assumption can be easily undermined, as we shall demonstrate in Section 4. In particular, it is far from obvious that Bayesianism is simpler, more unifying, or more rational than alternative frameworks given that advocates have yet to compare the relative non-empirical virtues of alternatives, nor have they systematically determined to what degree the Bayesian approach is more empirically adequate. We will have much more to say on this issue later on.

## 2.2 The Explanatory Power of Bayes

Biological cognitive systems access the world through their senses, which are viewed as sources of uncertain information about the state obtaining in the world at any given time. Describing the situation in statistical terminology, we may refer to the states obtaining in the world as 'environmental parameters', 'hidden states', or 'models', and the sensory information as 'sensory data' or 'evidence'. The problem faced by biological cognitive systems at any given time would thus consist in inferring which hidden state in the environment generated their sensory data. Unless this problem is effectively solved, reliable action, accurate perception, and adaptive learning would not be achievable.[2]

Often times, however, the values of environmental parameters are underdetermined by the available sensory data; that is, there are multiple, different states in the world that fit the sensory data received by a biological cognitive system at any given time. Because many different environmental states may fit the same piece of sensory data, processing sensory data alone is not sufficient to determine which state in the world caused it. Hence, sensory data underdetermine their environmental causes, which manifests a state of uncertainty within the system.

[2]There are many anti-Bayesian cognitive scientists who disagree with this idea, and hold instead that "cognition is not the representation of a pre-given world by a pre-given mind but is rather the enactment of a world and a mind on the basis of a history of the variety of actions that a being in the world performs" (Varela, Thompson, & Rosch, 1991, p. 9)

So, in part, the term ‘uncertainty’ in Bayesian cognitive science broadly characterizes this problem of underdetermination that biological cognitive systems must constantly solve.

To give an example, the sensory data generated by a convex object under normal lighting circumstances underdetermines its external cause (cf., Rescorla 2015a). There are at least two possible environmental states that can fit the available sensory data: the object in the world that caused the data is convex and the light illuminating the object comes from overhead; or the object is concave and the illuminating light comes from below. In order to perceive the world as being in one specific state, rather than as a superposition of two or more different states, cognitive systems must find some strategy to solve this problem of underdetermination.

Furthermore, uncertainty may also result from noise whose source can be internal or external to biological cognitive systems. In general, noise amounts to data received, but unwanted by a system. As a noisy signal contains more data than the original signal by itself, noise modifies the signal and extends the cognitive system’s freedom of choice in decoding it. This is an undesirable freedom to the extent that the adaptive behaviour the system can produce requires an appropriate degree of fidelity between original and decoded signals. Ultimately, noise poses a challenge for biological systems in estimating environmental parameters, and it is not at all an uncommon thing, for “noise permeates every level of the nervous system, from the perception of sensory signals to the generation of motor responses” (Faisal et al. 2008, p. 292).

Three sources of noise are characteristic of biological cognitive systems. The first source of noise lies in the thermodynamic or quantal transduction of the energy comprised by sensory signals into electrical signals. “For example, all forms of chemical sensing (including smell and gustation) are affected by thermodynamic noise because molecules arrive at the receptor at random rates owing to diffusion and because receptor proteins are limited in their ability to accurately count the number of signalling molecules” (Knill et al. 1996, p. 4). The second source of noise lies in biophysical features of ion channels, of synaptic transmission, of network interactions and random processes governing neural activations. These biophysical features introduce noise at the level of cellular signalling. A third source of noise lies in the transduction of signals carried by motor neurons into mechanical forces in muscle fibers. This transduction introduces noise in the signals supporting motor control, and can make motor behaviour highly variable even in the same types of circumstances when the same motor goal is pursued. In order to perform motor commands reliably and behave intelligently, biological systems must find some strategy to handle the noise introduced at different levels of neural processing.

We have now described how uncertainty tends to manifest within a biological cognitive system either from noise—random disturbances corrupting the sensory signals and processes of the system—or ambiguity—the underdetermination of percepts, as well as of other cognitive states, by sensory data. Whether caused by noise or surfacing from ambiguity, uncertainty goes hand in hand with variability in animals' environment and behaviour. On the one hand, the environment (including the body) constantly changes. On the other hand, even assuming that the environment does not change, cognitive systems' behaviour shows an ineliminable degree of variability. For example, if you reach for an object in the darkness, your visual and motor systems will lack relevant information about the location of the object. Your uncertainty about its location will be reflected in a lack of accuracy in any one reaching trial. If you try to reach for that object over and over again, you'll observe a large variability in your movement over reaching trials. Likewise, when a visual stimulus is hold as constant as possible, your visual perceptions of the stimulus will also vary over time. For biological systems to have accurate perceptions and to display reliable motor behaviour, they must find some way to tame this variability.

Biological cognitive systems would effectively deal with sensory and motor uncertainty if they were equipped with mechanisms that can solve the problem of underdetermination and mitigate the detrimental effects of noise. There is no doubt that Bayesianism provides a powerful framework for explaining the success of handling uncertainty. Within the theory, uncertainty is precisely specified by non-extreme probabilities, and the probabilities are updated upon obtaining new sensory information according to some rule, the most familiar of which is simple Bayesian conditionalisation. Bayesian biological systems would thus maintain, "at each stage of local computation, a representation of all possible values of the parameters being computed along with associated probabilities" (Knill and Pouget 2004, p. 713). With many advancements made in the Bayesian paradigm over recent decades, cognitive scientists are provided with a suite of algorithms and tools for precisely representing and computing the uncertainty of cognitive systems at any given time. These algorithms are currently employed in machine learning and statistics to solve problems of underdetermination and to mitigate detrimental effects of noise. Cognitive scientists may then choose the Bayesian framework to seek explanations of central aspects of cognition and behaviour.

## 2.3 Bayes and Uncertainty: A Natural Marriage?

The second step in the argument from uncertainty seeks to establish that Bayesianism *is* indeed the best for seeking explanations of how biological cognitive agents grapple with uncertainty. The focus is on non-empirical (or super-empirical) virtues of Bayes. By adopting the Bayesian approach, cognitive scientists would be able to explain most simply, most generally, and most rationally how biological cognitive agents solve the problem of underdetermination and handle the detrimental effects of internal noise. Hence, cognitive scientists *should* choose the Bayesian framework for discovering and assessing explanations of cognitive phenomena and behaviour. The idea that Bayes is in some sense the best for explaining how a system grapples with uncertainty is widely endorsed (e.g. Chater, Tenenbaum and Yuille 2006, p. 287; Doya et al. 2007; p. xi; Fiser et al. 2010, p. 120; Knill and Pouget 2004, p. 712; Maloney 2002, p. 145; Mamassian, Landy and Maloney 2002, p. 13; Orbán and Wolpert 2011, p. 1; Rescorla, 2015a, Sec. 2). Unfortunately, though, the idea is ambiguous, but hardly clarified.

On one account, the Bayesian framework provides a better language for representing uncertainty than competing alternatives. It therefore should be preferred as a way of modelling uncertainty and inference under uncertainty in cognitive science. On a different account, systems that implement Bayesian algorithms deal with uncertainty most effectively. Cognitive scientists should then prefer Bayesianism to alternatives, given that it is the best for understanding biological cognitive systems' adaptive behaviour in the face of uncertainty. While the first idea concerns the *representational* virtues of Bayesianism, the second idea concerns the *normative* character of the theory. Both sets of virtues concern non-empirical properties of Bayes.

As for the formal details, uncertainty is *represented* by probability within a Bayesian framework. Various rules of conditionalisation, such as simple Bayesian conditionalisation, Jeffrey's conditionalisation, or free-energy minimization, specify how belief is updated in the light of new information. Cognitive systems are assumed to entertain "beliefs" drawn from a hypothesis space $\mathcal{A}$. Beliefs concern what in the world could have caused the current sensory data E to the system. Each belief is associated with a prior probability $P(\mathrm{H})$, which represents the weight borne by the belief that H on the processes carried out by the system. At any given time, the system's beliefs satisfy the axioms of finitely additive probability. Probabilities are also assigned to (E, H) pairs, in the form of a *generative model* that specifies a joint probability distribution over sensory data and hypotheses about states in the world generating those data. Generative models represent likelihoods concerning how probable it is that the system would receive the current data E, given a hypothesized state H in the world, *viz.* $P(\mathrm{E}|\mathrm{H})$. With a generative model $P(\mathrm{E}|\mathrm{H})$, the current data E, and prior knowledge $P(\mathrm{H})$, the system computes the posterior conditional probability $P(\mathrm{H}|\mathrm{E})$, thereby reallocating probabilities across the hypothesis space in accordance with some learning rule.[3]

[3] Learning rules govern belief update, but they do not specify how the beliefs entertained by the system are used to produce a decision, action, or some other behavioural phenomenon. How the posterior is used to produce a decision requires the definition of a loss function, which

If such a concise framework provides us with a better way of *representing* inference under uncertainty than alternative frameworks, then the uncertainties a cognitive system face should be represented with a classical probability distribution, and the system's belief updates should be modelled by some Bayesian learning rule such as simple conditionalisation.

The idea that systems implementing Bayesian algorithms most effectively deal with uncertainty is often understood within David Marr's (1982) three levels of analysis framework (Griffiths, Vul, and Sanborn 2012). Marr's levels include the computational, the algorithmic, and the implementation level. The computational level specifies the problem to be solved in terms of some generic input-output mapping. In the case of Bayesian modelling in cognitive science, this is the problem of handling uncertainty. If the task is one of extracting some property of a noisy sensory stimulus, the generic input-output mapping that defines the computational problem is a function mapping the noisy sensory data to an estimate of the stimulus that caused that data. It is generic in that it does not specify any class of rules for generating the output. This class is defined at the algorithmic level. The algorithm specifies how the problem can be solved.

Bayesianism has traditionally been used as a "computational" benchmark of ideal performance in perceptual tasks (Geisler 2011), since it can define the problem agents are to solve and its "'optimal" solution (e.g., Griffiths et al. 2010; Griffiths et al. 2012). However, the success of Bayesian cognitive scientists in fitting many different sets of behavioural data has motivated uses of Bayesian models as process models, which specify algorithms and representations that cognitive systems can employ to actually solve a given problem (Colombo and Seriès 2012; Rescorla 2015a; 2015b; Zednik and Jaekel 2014). Further, more and more attention is being paid to how probabilistic representations and Bayesian algorithms can be implemented in neurally plausible mechanisms (Pouget et al. 2013; Ma and Jazayeri 2014).

In fact, Bayesian cognitive scientists are interested in all three Marr's levels of analysis (Griffiths, Vul and Sanborn 2012), and "the real interest [of Bayesian models] comes from the stronger notion that human beings might actually use the apparatus of probability theory to make their decisions, explicitly (if not consciously) representing prior probabilities, and updating their beliefs in an optimal, normatively sound fashion based on the mathematics of probability theory" (Marcus and Davis 2013, p. 2358).

To illustrate, let us consider the case of visual perception. Input data to the visual system consist of the light pattern that strikes the retina, initiating a cascade

specifies the relative cost of making a certain decision based on a certain belief. To determine the optimal decision available at a given time, the system needs to compute the estimated loss for any given decision and belief.

of chemical and electrical events that trigger neural signals. The beliefs (or hypotheses) entertained by the visual system consist of neurally encoded probability distributions of states in the world that could have produced that light pattern. Based solely on input data, the system cannot determine which particular state in the world caused the patterned excitation of the cones and rods in the retina. Any patch of retinal stimulation could correspond to an object of any size and almost any shape. However, if the system deploys "knowledge" about which size and shape are more probable a priori, it can determine which state would be most probable to produce the retinal input data. By applying a simple Bayesian rule of conditionalisation to combine prior knowledge with the likelihood of the state in the world producing the input data, the system identifies the worldly state with the *highest posterior probability*, which would provide an optimal solution to the problem of underdetermination.

Prior, probabilistic information embodied in neurons' receptive fields—*viz.* in the portion of sensory space that can elicit neural responses when stimulated—can be used and processed in a Bayesian fashion also to handle the effects of noise. The basic strategy is as follows: "If the structure of the signal and/or noise is known it can be used to distinguish signal from noise," which is essential to producing *accurate* perceptions and *reliable* motor behaviour (Faisal et al. 2008, p. 298). Neurons' prior probabilistic "knowledge" about the expected statistical structure of a signal from a given source of information would allow a biological system system to compensate for noise, and to give more weight to more reliable (less noisy) signals in its processes.

If Bayesianism defines a benchmark for optimal behaviour in the face of uncertainty, and provides us with algorithms that cognitive systems might implement in order to find optimal solutions to the problems they face, then Bayesianism should be preferred to seek explanations for cognitive systems' behaviour in the face of uncertainty.

# 3 Representing Uncertainty and Explaining Optimally

The previous section illustrated how one might come to think that the Bayesian framework is the best for explaining how cognitive systems deal with uncertainty due to its representational virtues or its normative virtues. Unfortunately, neither set of virtues provides cognitive scientists with a compelling reason to prefer the Bayesian framework over alternatives as one's basis for discovering and assessing explanations of cognitive phenomena. To understand why, let's ask two questions: What are the properties of Bayesianism that contribute to its representational power? What are the properties of the theory that contribute to its capacity to define computational problems and to provide “optimal” solutions to such problems?

## 3.1 Representing Uncertainty

Any given explanatory framework has certain epistemic, non-empirical properties that contribute to its representational power. Two such properties are its simplicity and unificatory power. If a framework $F$ possesses more of these properties or each of the properties to greater extent in comparison to an alternative framework $F^*$, then there is reason based on non-empirical virtues to prefer $F$ to $F^*$ as one's working framework for scientific explanation.

In the proceeding discussion, simplicity may plausibly be understood as a measure of the number and conciseness of the framework's basic principles while unification may be understood as a measure of the number of different kinds of phenomena to which the framework can be applied. Although the simplicity of a framework might bring with it pragmatic advantages like being more perspicuous and easier to use and to manipulate, unification can bring with it epistemic advantages related to explanation (Friedman 1974; Kitcher 1989) and to confirmation (Sober 2003).

Considering the Bayesian approach in cognitive science, the framework achieves simplicity. In particular, the basic principles of the theory are threefold: (i) an agent's beliefs that differ in strength are modeled by determinate probabilities; (ii) at any given time, an agent's beliefs obey the axioms of finitely additive probability; (iii) over time, an agent updates their beliefs according to a rule of conditionalisation. These three principles allow cognitive scientists to formulate research questions in a compact and precise manner (Chater et al. 2006, p. 287). Additionally, the language has great unifying power (e.g. Tenenbaum et al., 2011, p. 1285). In fact, it offers a common, encompassing, and flexible mathematical language for studying a wide variety of phenomena (Knill and Pouget 2004; Griffiths et al. 2010).

Despite the apparent virtues, many Bayesian models of real-world, high-dimensional tasks are hard to formulate and manipulate. The first challenge is choosing a suitable model and prior. A suitable model should not limit unduly the form of probability distributions (e.g., normal) or functions (e.g. linear), which are involved in the solution of a given cognitive task. A suitable prior should not rule out serious candidate hypotheses by assigning them zero probability, nor should it spread uniform mass over hypotheses. It should be noted that critics have pointed out that Bayesian cognitive scientists often select priors and likelihood functions post-hoc, in response to the data (Bowers & Davis 2012a). The second challenge concerns computing the posterior distribution, which is intractable for most real-world problems and calls for approximations and heuristics that might themselves be computationally intractable (Kwisthout, Wareham, and van Rooij 2011).

As for unification, although the theory has been used to fit an impressive range of data from a diverse variety of cognitive and behavioural tasks, this kind of unifying power does not obviously have explanatory or confirmatory import (Colombo and Hartmann 2015). Similar to Bayesianism, the language of Lagrangian field theory can be used for studying many kinds of systems. For example, it can be applied both to the behaviour of a system of gravitational masses and that of an electric circuit. This fact, however, does not warrant the conclusion that we thus have a common explanation of the behaviour of both systems (Woodward 2014, 5.4).

## 3.2 Explaining Optimally

Ignoring our latest contentions for the moment, Bayesianism undoubtedly appears to be an ideal framework to many cognitive scientists given that it sets an intuitive normative standard for how adaptive agents should combine and weigh different beliefs, how they should update their beliefs upon receiving novel information, and how they should make decision under uncertainty (Doya et al. 2007; Griffiths et al. 2010; see also Bovens and Hartmann 2003).

The normative force of the theory tends to depend on an agent's degrees of belief obeying the probability calculus and being updated via conditionalisation, which in turn is typically justified by appealing to (synchronic and diachronic) Dutch book arguments or to Cox's (1946) theorem. While Dutch book arguments purport to establish that it is practically irrational for an agent to have degrees of belief that violate the axioms of the probability calculus due to sure loss in expectation, Cox's (1946) theorem yields a more "commonsense" justification showing that any rational measure of belief is isomorphic to a probability measure.

Alternatively, justification for the rationality of Bayesianism may be grounded in considerations of accuracy rather than betting behaviour where the accuracy of a belief is determined by its "closeness to truth" (Joyce 1998). Accuracy-based justifications involve the use of a proper scoring rule for measuring the accuracy of a *belief function* at a world $w$ where a scoring rule gradually penalizes belief functions as the distance from the ideal belief function at a world $w$ increases, e.g. a belief $P(\mathrm{H}) = 0$ is given a maximum penalty $x$ if $w \in \mathrm{H}$ and a lesser penalty $y < x$ otherwise. The accuracy of a belief state is reflected in the total score given to one's beliefs over an *opinion set* (Pettigrew 2016). Recent results show that, for some scoring rule, Bayes agents will minimize the overall inaccuracy of their belief state synchronically and diachronically (e.g. Leitgeb and Pettigrew, 2010a,b).

Both kinds of justifications have been challenged, however. Douven (1999), for example, demonstrates that there are packages of decision and inference rules that include non-Bayesian update rules that are not subject to dynamic Dutch books. Even if non-Bayesian approaches to reasoning and decision making make agents vulnerable to dynamic Dutch books, non-Bayesian approaches have advantages that might outweigh the risk of suffering from a monetary loss. Douven (2013) showed that the beliefs of agents employing some non-Bayesian learning rule sometimes converge to the truth faster than those of their Bayesian rivals. Furthermore, accuracy arguments, which—unlike Dutch book arguments—aim to provide a direct justification of the epistemic rationality of Bayesianism, trade on ambiguity. There are different ways of understanding the epistemic goal of inaccuracy minimization. On some plausible understanding of *inaccuracy minimization*, Bayesianism is outperformed by alternative approaches (Douven 2013; Douven and Wenmackers 2015). In addition, there has been recent doubt casted on the viability of such an accuracy program (Littlejohn 2015).

# 4 Trafficking with Uncertainty: A Zoo of Approaches

Early on, we alluded to a significant problem that cognitive scientists who endorse the argument from uncertainty face, which concerns the neglect of different, yet promising, models that capture inference under uncertainty. The neglect has blinded them from recognizing the limitations of Bayesianism, especially since the literature surrounding alternative frameworks in computer science and statistics make such limitations transparent.

With this problem brought into focus, we contend that it should not be taken for granted that states of uncertainty involved in cognitive systems' processes are best represented and explained by Bayesianism, namely, because a systematic comparison of the non-empirical virtues of Bayes and other models remains to be seen within the cognitive sciences. In absence of such comparison, the Bayesian approach has unjustifiably gained its support by shielding itself from other plausible theoretical frameworks.

In this section, we correct the deficiency by detailing five different formal frameworks including Dempster-Shafer theory, imprecise probability, possibility theory, ranking theory, and quantum probability theory.[4] Of course, what follows is not intended to be an exhaustive literary review (see e.g. Halpern 2003; Huber 2014).

[4] While each framework provides a representation of uncertainty within the scope of probability theory, there are implicit, non-probabilistic approaches to uncertainty also (see e.g. Simoncelli 2009; Drugowitsch and Pouget, 2012). Although such theories are not discussed here, pointing them out at least expands the set of alternatives to the classical Bayesian framework. Taking into account a wider range of alternatives will make the case against the argument from uncertainty more convincing.

Nevertheless, the merits of each model considered suffice to undermine the assumption that Bayes supplies cognitive scientists with the *most simple*, *unifying*, and *rational approach* for explaining uncertainty-involving cognitive phenomena.

Toward the end of this section, we illustrate that the Bayesian approach is limited in handling a range of uncertainty-involving cognitive phenomena. To drive the point home, we turn our attention to a mathematically distinct framework of quantum probability and its application in cognitive science. Given that quantum probability is a genuine alternative theory to Bayesianism, the evidence supporting inference under uncertainty via quantum probability reduces support for Bayes, further undermining the assumption that the latter best represents and explains how cognitive systems handle uncertainty. Section 4.8 will then support our recommendation that the cognitive science community, if it cares about the health of their discipline and believe in its promise for accumulating knowledge about human mind, should not exhaust all of their efforts within a single approach with an argument from Bayesian confirmation theory.

## 4.1 The Dempster-Shafer Framework

The Dempster-Shafer (D-S) theory of evidence offers a more general framework than Bayesianism, for it represents "degrees of belief" through a pair of non-additive functions rather than a single additive probability function and accommodates learning through Dempster's rule for aggregating evidence instead of simple conditionalisation (Shafer 1992).

There are three functions used in modeling uncertainty in this framework: a *mass* function, a *belief* function, and a *plausibility* function. Let $W$ be a finite set of states. A mass function $m$ is a mapping of subsets from the power set or *frame of discernment* $\wp(W)$ (the set of all subsets including $\oslash$) to the unit interval [0,1], where $m(\oslash) = 0$ and the sum of the masses for all $\mathrm{X} \subseteq W$ is 1. The D-S belief function, Bel, and the plausibility function, Pl, define a lower and upper bound, respectively, representing the levels of *support* and *lack of evidence against* each element $\mathrm{X} \in \wp(W)$. The lower bound, Bel, is defined as the sum of the masses for all the subsets of a *set of interest*, $\mathrm{X} \in \wp(W)$. The upper bound, Pl, is defined as the sum of the masses for all the subsets that intersect the set of interest. Consider a simple example in Table 1 involving a finite set of states $W = \{\omega_1, \omega_2, \omega_3\}$.

We see that $\sum_{\mathrm{X} \subseteq W} m(\mathrm{X}) = \mathrm{Bel}(W) = \mathrm{Pl}(W) = 1$, thus entailing a complete lack of uncertainty with respect to the sure event $W$. For each proper subset of $W$, however, the same cannot be said. Levels of uncertainty associated with each proper subset are realized by Bel() and Pl(). Take the subset $\{\omega_1, \omega_3\}$, for example, which we will label X. The sum of the masses assigned to $\{\}$, $\{\omega_1\}$, $\{\omega_3\}$, and $\{\omega_1, \omega_3\}$ is 0.3, which is the lower bound, Bel(X). The sum of all the masses of subsets that intersect the set of interest is the upper bound, $\mathrm{Pl(X)} = 0.8$. The pair

| | $\{\}$ | $\{\omega_1\}$ | $\{\omega_2\}$ | $\{\omega_3\}$ | $\{\omega_1, \omega_2\}$ | $\{\omega_1, \omega_3\}$ | $\{\omega_2, \omega_3\}$ | $\{\omega_1, \omega_2, \omega_3\}$ |
|---|---|---|---|---|---|---|---|---|
| Mass | 0 | 0.1 | 0.2 | 0.1 | 0.1 | 0.1 | 0.3 | 0.1 |
| Belief | 0 | 0.1 | 0.2 | 0.1 | 0.4 | 0.3 | 0.6 | 1 |
| Plausibility | 0 | 0.4 | 0.7 | 0.6 | 0.9 | 0.8 | 0.9 | 1 |

Table 1: Dempster-Shafer Example

generates a "belief interval", $[0.3, 0.8]$, which in addition captures *partial ignorance*, illustrated by the difference between Bel and Pl with Bel $\leq$ Pl.

A further distinctive feature of D-S theory is that the union of disjoint events are believed at least as strongly as the sum of the lower bounds given for each event instead of there being strict equality between the probability of the union and the sum of the marginal probabilities. This implies that Bel is *superadditive*, i.e. $\mathrm{Bel}(\mathrm{X} \cup \mathrm{Y}) \geq \mathrm{Bel}(\mathrm{X}) + \mathrm{Bel}(\mathrm{Y})$ for all disjoint elements X, Y $\in \wp(W)$. The conjugate Pl, on the other hand, is *subadditive*, i.e. $\mathrm{Pl}(\mathrm{X} \cup \mathrm{Y}) \leq \mathrm{Pl}(\mathrm{X}) + \mathrm{Pl}(\mathrm{Y})$ for all disjoint elements X, Y $\in \wp(W)$. So, "degrees of belief" in D-S theory turn out to be non-additive, unlike in classical Bayesian theory.

In regard to belief updating, D-S theory employs Dempster's rule of combination for aggregating mass functions associated with information from multiple, independent sources:

$$(m_1 \otimes m_2)(X) = \frac{\sum_{\{X_1, X_2 : X_1 \cap X_2 = X\}} m_1(X_1) m_2(X_2)}{1 - K}$$

where $K = \sum_{\{X_1, X_2 : X_1 \cap X_2 = \oslash\}} m_1(X_1) m_2(X_2)$, $X \neq \oslash$, and $m(\oslash) = 0$. This rule corresponds to a normalized joint operation in which information is combined by favouring the agreement between the sources and ignoring all the conflicting evidence (Dempster 1968).

## 4.2 The Imprecise Probability Framework

In imprecise probability theory (IP), belief is usually modeled by a *credal set*, which is a non-empty set of probability functions $\mathcal{P}$ closed under convex combinations, and each probability function $P \in \mathcal{P}$ is defined on an algebra $\mathcal{A}$ generated by means of a finite set of states $W$ (see Levi 1974; Walley 1991; Joyce 2010).

Credal sets are bounded where the lower bound with respect to an event X is defined by a *lower probability* $\underline{\mathrm{P}}(\mathrm{X}) = \inf\{P(\mathrm{X}) : P \in \mathcal{P}\}$ for all X $\in \mathcal{A}$. The upper bound is defined by an *upper probability* $\overline{\mathrm{P}}(\mathrm{X}) = \sup\{P(\mathrm{X}) : P \in \mathcal{P}\}$. Lower and upper probabilities are conjugates such that $\underline{\mathrm{P}}(\mathrm{X}) = 1 - \overline{\mathrm{P}}(\mathrm{X}^c)$ and $\overline{\mathrm{P}}(\mathrm{X}) = 1 - \underline{\mathrm{P}}(\mathrm{X}^c)$. Given these conjugacy relations, it is only necessary to specify either a lower $\underline{\mathrm{P}}$ or an upper $\overline{\mathrm{P}}$ (the same goes for Bel and Pl in DS-theory) since they automatically define each other. As typical in the literature, IP proponents have established a preference for the language of (coherent) lower probabilities relative to a lower probability space $(W, \mathcal{A}, \mathcal{P}, \underline{\mathrm{P}})$ (Pederson and Wheeler 2014).[5]

Like in D-S theory, imprecise probability is distinct from classical probability in that the lower $\underline{\mathrm{P}}$ and upper $\overline{\mathrm{P}}$ functionals are non-additive. Instead, $\underline{\mathrm{P}}$ is superadditive and $\overline{\mathrm{P}}$ is subadditive: $\underline{\mathrm{P}}(\mathrm{X}\cup\mathrm{Y}) \geq \underline{\mathrm{P}}(\mathrm{X})+\underline{\mathrm{P}}(\mathrm{Y})$ and $\overline{\mathrm{P}}(\mathrm{X}\cup\mathrm{Y}) \leq \overline{\mathrm{P}}(\mathrm{X})+\overline{\mathrm{P}}(\mathrm{Y})$, for all X, Y $\in \mathcal{A}$. If we assume, as typically done, that $\mathcal{P}$ is closed under convex combinations—that is, for any $\lambda \in [0, 1]$ and $P_1, P_2 \in \mathcal{P}$, $\lambda P_1 + (1 - \lambda)P_2 \in \mathcal{P}$—then beliefs are represented by non-additive, interval-valued, imprecise probabilities: $\mathcal{P}(\mathrm{X}) = [\underline{\mathrm{P}}(\mathrm{X}), \overline{\mathrm{P}}(\mathrm{X})]$, for all X $\in \mathcal{A}$. On the surface, lower probability and belief functions look very much alike.

By contrast, though, IP is more general than Dempster-Shafer given that every D-S belief interval is an imprecise probability, but not every imprecise probability is a D-S belief interval (Huber 2016). Furthermore, IP collapses into Bayesianism with less restriction. To demonstrate, suppose that $\mathcal{P}$ is a singleton set, i.e. $\{P\}$, for all events in a respective algebra. Then, the probability function $P$ realizes the lower and upper probabilities, i.e. $\mathrm{P} = \underline{\mathrm{P}} = \overline{\mathrm{P}}$. What we learn from this fact is that $\mathcal{P}$ can in principle always be Bayesian, but whether Bel is Bayesian or not will ultimately depend on the masses assigned to subsets.

Additionally, IP more closely resembles Bayes in its belief updating method. Updating proceeds by way of conditioning each individually precise $P \in \mathcal{P}$ on new information E, assuming $P(\mathrm{E}) > 0$. The result is a set of conditional probabilities $\mathcal{P}(.|\mathrm{E})$ bounded by lower and upper conditional probabilities, $\underline{\mathrm{P}}(.|\mathrm{E})$ and $\overline{\mathrm{P}}(.|\mathrm{E})$. In the instance that $\mathcal{P}(.)$ and $\mathcal{P}(\mathrm{E})$ are singleton sets, $\mathcal{P}(.|\mathrm{E})$ is also a singleton set, i.e. $\mathcal{P}(.|\mathrm{E}) = \{P(.|\mathrm{E})\}$, yielding the same result as the classical model. Again, imprecise probability collapses into Bayesianism with less restriction since $\mathcal{P}(.|E)$ can in principle always be Bayesian whereas Dempster's rule is a more limiting case. We conclude that IP retains many of the benefits of Bayesianism, but it also has the expressive power that Dempster-Shafer enjoys.

## 4.3 The Possibility Framework

Possibility theory was originally inspired by ideas from fuzzy logic aiming at accommodating vagueness (Zadeh 1975). Using possibility theory for the purpose of measuring degrees of uncertainty rather than degrees of truth, a possibility measure, $\Pi$, models "the knowledge of an agent (about the actual state of affairs)

[5] Alternatively, Williams (2007) gives a more general theory of *previsions* taking the upper as primitive rather than the lower. For our discussion, nothing significant depends on whether we choose the lower or upper as primitive.

distinguishing what is plausible from what is less plausible, what is the normal course of things from what is not, what is surprising from what is expected" (Dubois and Prade 2007).

We begin the explication by defining a *possibility distribution* $\pi$ on a finite set of states $W$. The function $\pi$ maps a state $w \in W$ to a real number in the unit interval $[0, 1]$ and $\pi(w) = 1$ for at least one $w \in W$. From a possibility distribution, we can introduce a possibility measure $\Pi : \mathcal{A} \rightarrow \mathbb{R}$ that assigns 0 to $\oslash$ and 1 to $W$. For any $\mathrm{X} \in \mathcal{A}$, $\Pi(\mathrm{X}) = \sup_{w \in \mathrm{X}} \pi(w)$. Basically, $\Pi$ gives the degree to which an event is possible with 1 being maximally possible and 0 being impossible. A possibility measure also defines a conjugate necessity measure $\mathcal{N} : \mathcal{A} \rightarrow \mathbb{R}$ and $\mathcal{N}(\mathrm{X}) = \inf_{w \in \mathrm{X}} \pi(w)$ for all $\mathrm{X} \in \mathcal{A}$. A necessity measure $\mathcal{N}$ gives the degree to which an event is necessary.

Distinct from additive probability functions, a possibility function, $\Pi$, has the property of "maxitivity". Accordingly, if X and Y are disjoint sets, then $\Pi(\mathrm{X} \cup \mathrm{Y}) = \max(\Pi(\mathrm{X}), \Pi(\mathrm{Y}))$. This means that the union of disjoint sets is at least as possible as the maximally possible disjoint set, yet the union is no more possible than such set—hence, subadditivity. While $\Pi(\mathrm{X})$ is an upper bound with respect to uncertainty toward X, $\mathcal{N}(\mathrm{X})$ is the lower bound where $\mathcal{N}(\mathrm{X}) = 1 - \Pi(X^c)$. Consequently, we obtain a dual $\mathcal{N}(\mathrm{X} \cap Y) = \min(\mathcal{N}(\mathrm{X}), \mathcal{N}(\mathrm{Y}))$ (see Huber 2016).

In regard to conditional possibilities, if $\Pi(\mathrm{Y}) > 0$ and the set X is non-empty, then the canonical way to incorporate information is as follows:

$$\Pi(\mathrm{X}|\mathrm{Y}) = \begin{cases} 1 & \text{if } \Pi(\mathrm{X} \cap \mathrm{Y}) = \Pi(\mathrm{Y}); \\ \Pi(\mathrm{X} \cap \mathrm{Y}) & \text{if } \Pi(\mathrm{X} \cap \mathrm{Y}) < \Pi(\mathrm{Y}). \end{cases}$$

The difference between conditioning in possibility theory and Bayesian conditioning is that $\Pi(\mathrm{X} \cap \mathrm{Y})$ cannot be the product $\Pi(\mathrm{X}|Y)\Pi(\mathrm{Y})$ in an ordinal setting and so $\times$ is replaced by min.

With the bigger picture in mind here, $\Pi$ and $\mathcal{N}$ are similar to Pl and Bel, respectively. In fact, if a mass function $m$ on a finite frame of discernment is consonant by assigning positive mass only to an increasing sequence of sets, then a plausibility function Pl, relative to $m$, is a possibility measure (see Theorem 2.5.4 in Halpern 2003). With $\mathrm{Bel}_m$ being the conjugate of $\mathrm{Pl}_m$, it follows that $\mathrm{Bel}_m$ is a necessity measure. Thus, possibility theory is a limiting case of D-S and also IP.

## 4.4 The Ranking Framework

Ranking functions measure how surprising it would be if an event were to occur (or if some hypothesis is true). Formally, a ranking function $\mathcal{K} : 2^W \rightarrow \mathbb{N} \cup \{\infty\}$ models the degree of disbelief or surprise assigned to a subset of a finite space $W$ (Spohn 2012). We say that a subset X is surprising or disbelieved just in case its rank is positive, i.e. $\mathcal{K}(\mathrm{X}) > 0$. Subsets that are disbelieved are ranked gradually by the natural numbers to a max of $\infty$. The higher rank, the higher degree of surprise. Intuitively, the empty set should be disbelieved to the highest degree.

On the other hand, if any subset X of $W$ is not at all surprising, then the subset is assigned a rank of 0. Intuitively, $W$ should never be disbelieved for it is not surprising that one of the states in $W$ will obtain. However, unsurprisingness does not necessarily imply that if X is assigned a rank of 0, then X is believed. The subset X is said to be believed just in case its complement is disbelieved, i.e. $\mathcal{K}(\mathrm{X}^c) > 0$. Otherwise, a rank of 0 assigned to $X$ and $X^c$ would seem to suggest suspension of judgment since one has yet to come to disbelieve one of the disjoint sets.

Moreover, in reverse of possibility theory, a union of (non-empty) sets, X and Y, has a rank equal to the rank of the lowest ranking set. So we observe the connection to possibility measures with 0 and 1 being replaced by $\infty$ and 0 and max replaced with min. As for conditional rankings, they are defined as such: $\mathcal{K}(\mathrm{X}|\mathrm{Y}) = \mathcal{K}(\mathrm{X} \cap \mathrm{Y}) - \mathcal{K}(\mathrm{Y})$. Using conditional ranks, the main rules for updating on new information correspond to Bayesian conditionalisation with an analogue to Bayes' rule: $\mathcal{K}(\mathrm{X}|\mathrm{Y}) = \mathcal{K}(\mathrm{Y}|\mathrm{X}) + \mathcal{K}(\mathrm{X}) - \mathcal{K}(\mathrm{Y})$.

## 4.5 The Quantum Probability Framework

Quantum probability theory is a geometric model of uncertainty. It uses fragments of the language of mathematical probability, but outcomes are distinctively represented as subspaces of varying dimensionality in a multidimensional Hilbert space, which is a vector space used to represent all possible outcomes for questions asked about a system. Unit vectors correspond to possible states of the system, and embody knowledge about the states of the system under consideration.

Probabilities of outcomes are determined by projecting the state vector onto different sub-spaces and computing the squared length of the projection. The determination of probabilities is context- and order-dependent, as individual states can be superposition states and composite systems can be entangled. Thus, while in the Bayesian framework $P(\mathrm{X} \cap \mathrm{Y}) = P(\mathrm{Y} \cap \mathrm{X})$, the commutative property in quantum probability does not always hold. More generally, unlike in Bayesianism, quantum probability does not obey the law of total probability (see Rédei and Summers 2007 for an introduction to the theory).

Incompatibility in quantum probability theory entails that it is impossible to concurrently assign a truth-value to two hypotheses. Psychologically, two incompatible hypotheses in this sense can be processed only serially because the processing of one hypothesis interferes with the other. Given hypotheses A and B, for example, if A is true at a certain time, then B can be neither true nor false at that time. Conjunctions between incompatible hypotheses are then defined in a sequential way as “A and then B.”

One advantage of using the quantum probability framework is that it allows for explanations of cognitive systems being in a superposition of different states. Superposition can give rise to a specific kind of uncertainty that is dependent on the fuzziness and ambiguity of information and that characterizes the ambivalence of many of our everyday judgments. Additionally, entanglement tracks the interdependencies between different parts of complex cognitive systems. In entangled systems, it is not possible to define a joint probability distribution from the probability distributions of variables corresponding to different constituent parts. In such systems, changes in one constituent part entails instantaneous changes in another part.

Courtesy of interference, superposition and entanglement, we are able to explain the conjunction fallacy, non-compositional conceptual semantics, order effects in perception, and violations of the sure thing principle (Busemeyer and Bruza 2012).

## 4.6 Heir to the Throne?

Each framework sketched has some advantage over Bayesianism, but limitations also. The Dempster-Shafer theory, for instance, has an advantage of representing states of complete ignorance without precise degrees of belief: 0 mass everywhere except for the sure event.[6] Furthermore, combining evidence with Dempster’s rule has the desirability of relaxing strong independence assumptions. Upon gathering new evidence, beliefs should be determined by combining the vacuous belief function with the total evidence. Finally, beliefs and evidence are represented by the same type of mathematical objects, *viz.* belief functions. All of this suggests that Bayes might not be the most *unifying* or explanatory theory after all.

A problem for D-S theory, however, is that inference is computationally inefficient compared to Bayesian inference. The inefficiency stems from evidence being represented by a belief function that is induced by a mass function on the frame of discernment instead of a probability distribution over a partition. Combining evidence by Dempster’s rule increases computational complexity as the number of possible states increases. In an attempt alleviate the complexity issue, though, Shafer and Tversky (1985) emphasised that “[t]he usefulness of one of these formal languages [i.e., the Bayesian and the D-S language] for a specific problem may depend both on the problem and on the skill of the user... A person may find one language better for one problem and another language better for another” (Shafer and Tversky 1985, p. 311).

[6]See Norton (2007, 2011) for a series of challenges that the Bayesian faces in representing complete ignorance. Norton, however, is not convinced that Dempster-Shafer theory does any better, but we will leave those matters to the side in this paper.

Moreover, one should recognize that D-S and Bayes are not incompatible theories, despite their dissimilarities. Some precise probability distributions are special cases of D-S belief functions and some Bayesian conditional probability distributions are special instances of applying Dempster's rule. D-S theory, then, is a generalization of Bayesianism. Pitting the theories against each other makes little sense from a confirmation-theoretic standpoint since any evidence confirming Bayes confirms D-S theory also.

Like Dempster-Shafer theory, imprecise probability is a generalization of Bayes. So again, any evidence that confirms Bayes also confirms IP. This logical fact does not provide reason to ignore the more expressive IP framework. In fact, there are explanatory reason for why Bayesian cognitive scientists should be compelled to adopt the broader language of imprecise probability to capture several cognitive phenomena. For instance, Daniel Ellsberg (1961) pointed out the phenomenon of *ambiguity aversion* and demonstrated that the preferences of tested subjects are inconsistent within a Bayesian framework. The results have since been replicated, indicating that many are averse to ambiguity in decision making (see Camerer and Weber 1992 for a survey). Mathematical economists later joined imprecise probability together with a maxmin decision criterion, which plausibly explained the results of Ellsberg's experiments under an axiomatized scheme (Gilboa and Schmeidler 1989). This example from psychology shows that Bayes is neither the most *unifying* nor *rational* framework.

IP theory, however, has a computational inefficiency problem, similar to D-S theory, involving great computational complexity when updating convex sets of probabilities. The theory also encounters trouble when it comes to updating "trivial states of uncertainty" or the non-informative prior, $[0, 1]$, which has prompted some to rule out the vacuous prior (see Walley 1991: Rinard 2013). But in doing so, the theory becomes restricted to representing partial ignorance and loses the capability of representing complete ignorance. In order to recover such representation, IP theory will need to be extended with a non-Bayes updating rule. Until then, IP is limited in its modelling capabilities with respect to uncertainty-involving phenomena.

The possibility framework has a computational advantage over probability as "maxitivity" makes possibility measures compositional—*viz.* $\Pi(\mathrm{X} \cup \mathrm{Y})$ is determined by the maximum of $\Pi(\mathrm{X})$ and $\Pi(\mathrm{Y})$. Minimal computation indicates that possibility theory is at least *simpler* than Bayesianism. Within the larger picture, there are similarities between possibility theory and Dempster-Shafer theory in which a Pl function can be a possibility measure. However, possibility need not be restricted to a D-S interpretation. In general, possibility theory "can be seen either as a coarse, non-numerical version of probability theory, or as a framework for reasoning with extreme probabilities, or yet as a simple approach to reasoning with imprecise probabilities" (Dubois and Prade 2007). The upshot of possibility theory is its usefulness in assessing vague statements like 'Bob is tall' or 'the shirt is blueish'. Given its application to vagueness, possibility theory may offer cognitive scientists a more *unified* method for explaining reasoning under uncertainty with vague concepts. But since fuzzy approaches to uncertainty such as possibility theory are not isomorphic to probability theory, it could be suggested that Cox's theorem rules out possibility theory as a *rational* means of quantifying uncertainty (Lindley 1982; but see Colyvan 2008).

Ranking theory has parallels to possibility theory. But distinct from Bayesianism and the considered theoretical frameworks, proponents point out that ranking theory accommodates the everyday, categorical notion of belief (and disbelief), not just quantitative degrees of belief. On these grounds, they claim that the ranking theoretic approach has advantages over probabilistic approaches because it allows for everything that we can do with quantitative measures and also to tackle traditional problems in epistemology that center around the traditional tripartite concept of belief (Spohn 2012).[7] Ranking theory, then, can be thought of as more *unifying* than Bayesianism. In the cognitive sciences, ranking theory has received some attention especially in the AI community (e.g., Kern-Isberner and Eichorn 2014). However, its applications in experimental psychology are currently limited, since it is not obvious how to derive experimentally distinguishable predictions from ranking theory (but see Skovgaard-Olsen 2015).

Finally, quantum probability theory is uniquely based on a set of axioms allowing for an agent to be Dutch-booked. As we noted, Dutch book arguments do not provide a decisive reason for the superiority of Bayesianism for there are other ways of vindicating Bayes. Moreover, while quantum probability theory "is perhaps a framework for bounded rationality and not as rational as in principle possible" (Pothos and Busemeyer 2014, p. 2), courtesy of its unique properties, including superposition, entanglement, incompatibility, and interference, it accommodates empirical results related to order and context effects that are not plausibly captured within a Bayesian framework (Pothos and Busemeyer 2013). Such capability indicates that quantum probability is more *unifying*. An example of this will be detailed in the following subsection, which we turn to now.

## 4.7 Casting Further Doubt: Perceptual order Effects

The observations made in this section thus far undermine the assumption that Bayesianism is the best for representing and explaining uncertainty, and that it is

[7] Ranking theory is not the only approach to unifying degrees of belief and categorical belief. For example, Hannes Leitgeb's (2014) *stability theory of belief* proposes a unified theory also. Another approach, which sidesteps deductive closure of full belief, is Easwaran and Fitelson's (2015) accuracy-based account.

the most rational approach, too. However, virtually all Bayesian cognitive scientists have neglected the above possibilities in their work and proceeded by taking for granted the superiority of Bayesianism, or assuming that Bayesianism is the only game in town. This is ultimately a methodological problem hindering progress in cognitive science. In fact, recent findings in perception, judgment, and decision making have been shown to be unaccommodated by the Bayesian framework, but are plausibly explained by quantum probability.

For example, the experimental results of Conte and colleagues (2009) who investigated quantum-like interference effects in the perceptual domain show the failure of Bayesianism. In their experiment, participants were presented with ambiguous images, which could be perceived in two mutually exclusive manners. One group of participants was presented with a single image A and asked to make a binary choice between A = a or A = q on the basis of the manner in which they perceived the image at the instance of observation. Another group of participants was presented with two ambiguous images, B and A. After each presentation, participations had to make a binary choice between B = b and B = r, and between A = a and A = q.

By the law of total probability, Bayesian probability predicted the probability that a participant chooses A = a in any of the trials is: $P(\mathrm{A} = \mathrm{a}) = P(\mathrm{B} = \mathrm{b})P(\mathrm{A} = \mathrm{a} \mid \mathrm{B} = \mathrm{b}) + P(\mathrm{B} = \mathrm{r})P(\mathrm{A} = \mathrm{a} \mid \mathrm{B} = \mathrm{r})$. The results were inconsistent with this prediction.[8] They were instead consistent with the quantum probability prediction that participants' choices would be affected by quantum like interference where the context generated by making the first perceptual choice interfered with the second so that the participants' choices showed order effects implying non-commutativity. Since such interference effects are ubiquitous in psychology (e.g., Kvam, Pleskac, Yu, and Busemeyer 2015), but incompatible with natural predictions of Bayesianism, the quantum probability approach can account for some phenomena that cannot be captured by Bayes.

Although the evidence does not fully vindicate quantum probability, it does undermine the view that Bayesianism is the most empirically adequate framework. If the argument from uncertainty is to succeed in justifying the Bayesian approach in cognitive science, then other plausible theoretical frameworks should not be ignored. Unless the relative epistemic, empirical and non-empirical virtues of the Bayesian framework are probed against the virtues and limitations of alternatives on an array of case-studies, the choice may be too premature to focus scientific resources on this one approach. For it needs to be established that Bayesianism is actually the "most simple", "most unifying", and "most rational" framework to

[8] Bayesians will be quick to point out that context effects can straightforwardly be handled in Bayesian models by adjusting the likelihoods so that they no longer assume independent and identically distributed samples. While this move underwrites the flexibility of Bayes, it also highlights the risk of ad-hocery that may come with this flexibility

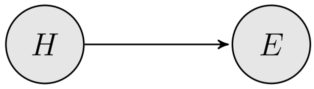

Figure 1: The Bayesian Network with the variables $H$ and $E$.

understand uncertainty-involving phenomena produced by cognitive systems.

## 4.8 A Bayesian Argument

In closing this section, we would like to give a Bayesian argument to the effect that the cognitive science, as well as the philosophical communities should not endorse Bayes too quickly. To do so, we consider a theory H which accounts for the evidence E. Introducing binary propositional variables $H$ and $E$ with the values H ("The hypothesis is true") and $\neg H$ ("The hypothesis is false') and E ("The evidence obtains") and $\neg E$ ("The evidence does not obtain"), the Bayesian Network depicted in Figure 1 describes the probabilistic relation between $H$ and $E$.[9]

We assume that H entails E, i.e. that the evidence is a deductive consequence of the theory. Hence, we set

$$P(\mathrm{H}) = h \quad , \quad P(\mathrm{E}|\mathrm{H}) = 1 \quad , \quad P(\mathrm{E}|\neg\mathrm{H}) = \alpha \tag{1}$$

with $0 < \alpha < 1$. It is then easy to show that the posterior probability of H, i.e. the probability of H after learning that E is true, is given by[10]

$$P(\mathrm{H}|\mathrm{E}) = \frac{h}{h + \alpha\bar{h}} > h \tag{2}$$

Hence, E confirms H. But how much? As one sees from eq. (2), $P(\mathrm{H}|\mathrm{E})$ is a decreasing function of $\alpha$. For $\alpha \to 0$, i.e. if we consider it to be impossible that an alternative theory (which is contained in the "catch-all" $\neg$H) accounts for the evidence, then $P(\mathrm{H}|\mathrm{E}) \to 1$. For $\alpha \to 1$, i.e. if we are convinced that an alternative theory accounts for the evidence, then $P(\mathrm{H}|\mathrm{E}) \to h$. Hence, if we consider it quite likely that an alternative theory accounts for the evidence, i.e. if we set $\alpha \approx 1$, then $P(\mathrm{H}|\mathrm{E}) \approx h$ and we won't get much confirmation for H after observing E.

This situation changes is we consider several independent pieces of evidence $E_1, \ldots, E_n$. Assuming $E_i \perp\!\!\!\perp E_j | H$ for $i \neq j = 1, \ldots, n$ and setting $P(\mathrm{H}) = h, P(\mathrm{E_i}|\mathrm{H}) = 1$ and $P(\mathrm{E_i}|\neg\mathrm{H}) = \alpha$ for $i = 1, \ldots, n$, we obtain

$$P(\mathrm{H}|\mathrm{E_1}, \ldots, \mathrm{E_n}) = \frac{h}{h + \alpha_{eff}\,\bar{h}} \,. \tag{3}$$

[9] See Bovens and Hartmann (2003) for an introduction to the application of Bayesian Network methods in epistemology and philosophy of science.

[10] Here and throughout we use the abbreviation $\bar{x} := 1 - x$.

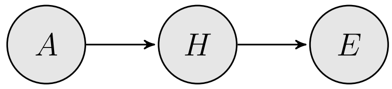


Figure 2: The Bayesian Network with the variables $H, E$ and $A$.

For large $n$, $\alpha_{eff} := \alpha^n \approx 0$ and hence $P(\mathrm{H}|\mathrm{E}) \approx 1$. Given that Bayesian Cognitive Science accounts for many different phenomena, this seems to justify taking Bayesian Cognitive Science very seriously. Note, however, that we made two important assumptions. First, we assumed that the different pieces of evidence $E_1, \ldots, E_n$ are independent (given H). This is controversial and needs to be justified on a case by case basis. Second, we assumed that all pieces of evidence are a deductive consequence of H, i.e. we assumed that $P(\mathrm{E_i}|\mathrm{H}) = 1$ for all $i = 1, \ldots, n$. This is controversial as H may make $\mathrm{E}_i$ only highly likely, and so we should set $P(\mathrm{E_i}|\mathrm{H})$ to a value smaller than 1. Assigning a value smaller than 1 to $P(\mathrm{E_i}|\mathrm{H})$ is also supported by the observation that $\mathrm{E}_i$ typically does not follow from H alone, but from H and some additional auxiliary assumptions. (This is the famous Duhem Quine Problem.) Hence, $P(\mathrm{E_i}|\mathrm{H}) < 1$ (for all $i = 1, \ldots, n$) which effectively lowers the posterior probability $P(\mathrm{H}|\mathrm{E_1}, \ldots, \mathrm{E_n})$.

To accept H, we would also like to make sure that the posterior probability of H is fairly high. As eq. (2) shows, the value of $P(\mathrm{H}|\mathrm{E})$ also depends on the prior probability of H (i.e. on $h$) and neglecting it would mean to commit the base-rate fallacy.

So let us now explore what we can say about the prior probability of H. We will argue that it depends on our beliefs about the existence of alternative theories that explain the evidence. To proceed with our analysis, we additionally introduce the binary propositional variable $A$ with the values A := "There is an alternative explanation for E" and ¬A accordingly and study the Bayesian Network depicted in Figure 2. We set

$$P(\mathrm{A}) = a \quad , \quad P(\mathrm{H}|\mathrm{A}) = \beta \tag{4}$$

with $0 < a, \beta < 1$. $\beta$ will be large if we believe that H is part of a better explanation (provided there is one) or if we believe that there can be multiple equally acceptable explanations for E. $\beta$ will be small if we believe that an alternative explanation will be better and eventually replace H. $\beta$ will also be small if one beliefs that either H or some alternative is right and that there can be only one explanation for E. Hence, if there is an alternative, this alternative might well be the true theory and hence one assigns a small value to $\beta$. Given that there are several more or less unexplored alternative theoretical frameworks in cognitive science (as argued above), it seems rational to assign a rather low value to the parameter $\beta$.

We also set

$$P(\mathrm{H}|\neg\mathrm{A}) = 1 \tag{5}$$

as there must be (or so we assume) an explanation for E. If A is false and there is no alternative explanation for E, then H has to be true.

With this, we calculate

$$P(\mathrm{H}) = h = a\beta + \bar{a}. \tag{6}$$

Hence, if one has good reasons to believe that $a$ is fairly large (i.e. if, as in the case of Bayesian Cognitive Science, alternatives to H are known and we can assume that they provide alternative explanations of E) and if $\beta$ is fairly small (as we argued), then the "prior" $h$ is relatively small and hence the posterior probability $P(\mathrm{H}|\mathrm{E})$ is relatively small.

We have therefore given a Bayesian argument to the effect that we have to be careful and not accept Bayesian Cognitive Science too quickly. It might well be worth the effort to systematically explore alternative theoretical frameworks and identify phenomena that can distinguish between the different explanatory theories.

# 5 Conclusion. Against Monopoly

If there is good reason to doubt that the Bayesian approach provides us with the best explanations for many cognitive phenomena, then there is good reason to remain agnostic about the truth of Bayesian models of cognitive phenomena and behaviour, contrary to what has been claimed in the philosophical literature (e.g., Rescorla 2015a; 2015b), and in the cognitive science literature (e.g., Knill and Pouget 2004; Ma et al. 2006; Friston 2009).

Facts about the institutional organization of contemporary scientific inquiry bolster this agnosticism, providing us with some explanation of why alternatives to Bayesianism have been neglected. As pointed out by Stanford (2015), the institutional apparatus of contemporary scientific inquiry has "served to reduce not only the incentives but also the freedom scientists have to pursue research that challenges existing theoretical orthodoxy or seeks to develop fundamental theoretical innovation." While this conservatism has fostered specialization in the sciences, it has also shielded popular theories and frameworks from comparison with relevant, under-considered alternatives. This conservatism, coupled with the neglect for available alternatives, pose a challenge to a realistic attitude towards Bayesian models of cognition. More important, this conservatism substantiates the concern that valuable scientific resources are being unjustifiably monopolized by Bayesians.

Currently, there is little doubt that Bayesianism is the most popular approach to represent and deal with uncertainty (Halpern 2003, p. 4). The tools which a Bayesian cognitive scientist can currently use to address problems of uncertain inference are more sophisticated than alternatives, being routinely used in neighboring fields like machine learning, artificial intelligence, and statistics. In comparison to Dempster-Shafer theory, possibility theory, ranking theory, and quantum probability theory, the Bayesian approach is more widespread in each of a wide variety of fields ranging from statistics to machine learning and AI (Poirier 2006; Jordan and Mitchell 2015). And the popularity of Bayesian modelling has been growing in cognitive science too, as evidenced by an increase in the number of articles, conference papers and workshops dedicated to Bayesian modelling of cognition and its foundations (cf. Kwisthout et al. 2011, note 1).

Given this popularity, and given that it is not obvious that the Bayesian framework enjoys special epistemic virtues in comparison to alternatives, many cognitive scientists' choice to carry out their research within the Bayesian framework can be plausibly explained in terms sociological factors connected with the reward structure of scientific institutions, which is biased towards conservatism (Stanford 2015). These sociological factors may have led more and more scientists to approach research questions within the Bayesian framework, while neglecting some of the alternatives. As more and more cognitive scientists have addressed research questions within the Bayesian framework, a division of cognitive labour has been fostered in the field. Sophisticated tools have been developed (Jordan and Mitchell 2015). Such tools have been exploited to approach problems at a higher level of specialization in both machine learning and human cognition (Gershman, Horvitz and Tenenbaum 2015). But if a higher degree of specialization arises within a scientific community with an incentive structure that strongly favours conservatism, then exploring and developing novel theoretical frameworks will happen with much more difficulty. As this will impact the trajectory of cognitive science, we believe – like Gigerenzer (1991) before us – that it is important to take a step back and evaluate whether the net result is the best way to advancing our understanding of how minds work.

# References

[1] Bovens, L. & Hartmann, S. (2003). *Bayesian Epistemology.* Oxford: Oxford University Press.

[2] Bowers, J.S. & Davis, C.J. (2012a). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin, 138,* 389-414.

[3] Bowers, J.S. & Davis, C.J. (2012b). Is that what Bayesians believe? Reply to Griffiths, Chater, Norris, and Pouget (2012). *Psychological Bulletin, 138,* 423-426.

[4] Busemeyer, J. R. & Bruza, P.D. (2012). *Quantum Models of Cognition and Decision.* Cambridge: Cambridge University Press.

[5] Camerer, C. & Weber, M. (1992). Recent developments in modeling preferences: uncertainty and ambiguity. *Journal of Risk and Uncertainty, 5,* 325-370.

[6] Chater, N., Goodman, N. D., Griffiths, T. L., Kemp, C., Oaksford, M., & Tenenbaum, J.B. (2011). The imaginary fundamentalists: the unshocking truth about Bayesian cognitive science. *Behavioral and Brain Sciences, 34,* 194-196.

[7] Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: conceptual foundations. *Trends in Cognitive Sciences, 10(7),* 287-291.

[8] Colombo, M. & Hartmann, S. (2015). Bayesian cognitive science, unification, and explanation. *The British Journal for the Philosophy of Science,* axv036.

[9] Colombo, M. & Seriès, P. (2012). Bayes in the brain–on Bayesian modelling in neuroscience. *The British Journal for the Philosophy of Science, 63(3),* 697-723.

[10] Colyvan, M. (2008). Is probability the only coherent approach to uncertainty? *Risk Analysis, 28,* 645-52.

[11] Conte, E., Khrennikov, A. Y., Todarello, O., Federici, A., Mendolicchio, L., & Zbilut, J. P. (2009). Mental states follow quantum mechanics during perception and cognition of ambiguous figures. *Open Systems and Information Dynamics, 16,* 1-17.

[12] Cox, R.T. (1946). Probability, frequency, and reasonable expectation. *American Journal of Physics, 14,* 1-13.

[13] Diaconis, P. & Zabell, S. (1982). Updating subjective probability. *Journal of the American Statistical Association, 77,* 822-830.

[14] Doya, K., Ishii, S., Pouget, A., & Rao, R.P.N. (2007). *Bayesian Brain: Probabilistic Approaches to Neural Coding.* Cambridge: MIT Press.

[15] Dempster, A. (1968). A generalization of bayesian inference. *Journal of the Royal Statistical Society, 30,* 205-247.

[16] Douven, I. (1999). Inference to the best explanation made coherent. *Philosophy of Science, 66,* S424-S435.

[17] Douven, I. (2013). Inference to the best explanation, Dutch books, and inaccuracy minimisation. *The Philosophical Quarterly, 63,* 428-444.

[18] Douven, I. & Wenmackers, S. (2015). Inference to the best explanation versus Bayes' rule in a social setting. *The British Journal for the Philosophy of Science,* axv025.

[19] Drugowitsch, J. & Pouget, A. (2012). Probabilistic vs. non-probabilistic approaches to the neurobiology of perceptual decision-making. *Current Opinion in Neurobiology, 22,* 963-969.

[20] Dubois, D. & Prade, H. (2007). Possibility theory. *Scholarpedia, 2,* 2074.

[21] Easwaran, K. & Fitelson, B. (2015). Accuracy, coherence, and evidence. *Oxford Studies in Epistemology, 5,* 61-96.

[22] Eberhardt, F. & Danks, D. (2011). Confirmation in the cognitive sciences: the problematic case of Bayesian models. *Mind and Machines, 21,* 389-410.

[23] Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics, 75,* 643-669.

[24] Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience, 9,* 292-303.

[25] Fiser, J., Berkes, P., Orban, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences, 14,* 119-130.

[26] Friedman, M. (1974). Explanation and scientific understanding. *The Journal of Philosophy, 71,* 5-19.

[27] Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences, 13,* 293-301.

[28] Geisler, W.S. (2011). Contributions of ideal observer theory to vision research. *Vision Research, 51,* 771-781.

[29] Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: a converging paradigm for intelligence in brains, minds, and machines. *Science, 349,* 273-278.

[30] Ghahramani, Z. (2015). Probabilistic machine learning and artificial intelligence. *Nature, 521,* 452-459.

[31] Gigerenzer, G. (1991). From tools to theories: a heuristic of discovery in cognitive psychology. *Psychological Review, 98,* 254-267.

[32] Gilboa, I. & Schmeidler, D. (1989). Maximin expected utility with non-unique prior. *Journal of Mathematical Economics, 18,* 141-153.

[33] Goldszmidt, M. & Pearl, J. (1996). Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence, 84,* 57-112.

[34] Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: causal maps and Bayes nets. *Psychological Review, 111,* 3.

[35] Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the Bayesians got their beliefs (and what those beliefs actually are): comments on Bower and Davis (2012). *Psychological Bulletin, 138,* 415-422.

[36] Griffiths, T. L., Vul, E., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science, 21,* 263-268.

[37] Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: exploring representations and inductive biases. *Trends in Cognitive Sciences, 14,* 357-364.

[38] Halpern, J. (2003). *Reasoning About Uncertainty.* Cambridge: MIT Press.

[39] Hartmann, S. & Rafiee Rad, S. (2016). Learning conditionals. MCMP Working Papers. LMU Munich.

[40] Huber, F. (2014). Formal representations of belief. *The Stanford Encyclopedia of Philosophy,* Edward N. Zalta (ed.), <http://plato.stanford.edu/archives/spr2014/entries/formal-belief/>.

[41] Jones, M.I. & Love, B.C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences, 34,* 169-188.

[42] Jones, M.I. & Mitchell, T.M. (2015). Machine learning: trends, perspectives, and prospects. *Science, 349,* 255-260.

[43] Joyce, J. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science, 65,* 573-603.

[44] Joyce, J. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives, 24,* 281-323.

[45] Kern-Isberner, G. & Eichhorn, C. (2014). Structural inference from conditional knowledge bases. *Studia Logica, 102,* 751-769.

[46] Kvam, P. D., Pleskac, T. J., Yu, S., & Busemeyer, J. R. (2015). Interference effects of choice on confidence: quantum characteristics of evidence accumulation. *Proceedings of the National Academy of Sciences, 112,* 10645-10650.

[47] Kersten, D., Mamassian, P., & Yuille, A. (2005). Object perception as Bayesian inference. *Annual Review of Psychology 55,* 271-304.

[48] Kitcher, P. (1989). Explanatory, unification, and the causal structure of the world. in P. Kitcher and W. Salmon (eds.), *Scientific Explanation. Minnesota Studies in the Philosophy of Science, Volume XIII,* 410-505.

[49] Knill, D.C. & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences, 27,* 712-719.

[50] Knill, D.C. & Richards, W. (1996). *Perception as Bayesian Inference.* Cambridge: Cambridge University Press.

[51] Körding, K. (2007). Decision theory: what "should" the nervous system do? *Science, 318,* 606-610.

[52] Kwisthout, J., Wareham, T., & van Rooij, I. (2011). Bayesian intractability is not an ailment that approximation can cure. *Cognitive Science, 35,* 779-784.

[53] Leitgeb, H. (2014). The stability theory of belief. *The Philosophical Review, 123,* 131-171.

[54] Leitgeb, H. & Pettigrew, R. (2010a). An objective justification of Bayesianism I: measuring inaccuracy. *Philosophy of Science, 77,* 201-235.

[55] Leitgeb, H. & Pettigrew, R. (2010b). An objective justification of Bayesianism II: the consequences of minimizing inaccuracy. *Philosophy of Science, 77,* 236-272.

[56] Levi, I. (1974). On indeterminate probabilities. *The Journal of Philosophy, 71,* 391-418.

[57] Lindley, D.V. (1982). Scoring rules and the inevitability of probability. *International Statistical Review, 50,* 1-26.

[58] Littlejohn, C. (2015). Who cares what you accurately believe? *Philosophical Perspectives, 29,* 217-248.

[59] Ma, W. J. & Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience, 37,* 205-220.

[60] Ma, W.J., Beck, J.M., Latham, P.E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience, 9,* 1432-1438.

[61] Mamassian, P., Landy, M. S., & Maloney, L. T. (2002). Bayesian modeling of visual perception. in R. Rao, M. Lewicki,, and B. Olshausen (eds.), *Probabilistic Models of the Brain: Perception and Neural Function.* Cambridge: MIT Press, 13-36.

[62] Maloney, L.T. (2002). Statistical decision theory and biological vision. in D. Heyer and R. Mausfeld (eds.), *Perception and the Physical World: Psychological and Philosophical Issues in Perception.* New York: Wiley, 145-189.

[63] Marcus, G.F. & Davis, E. (2013). How robust are probabilistic models of higher-level cognition? *Psychological Science, 24,* 2351-2360.

[64] Marr, D. (1982). *Vision.* New York: W.H. Freeman and Co.

[65] Norton, J.D. (2007). Disbelief as the dual of belief. *International Studies in the Philosophy of Science, 21,* 231-252.

[66] Norton, J.D. (2011). Challenges to Bayesian confirmation theory. in P.S. Bandyopadhyay and M. Forster (eds.), *Handbook of the Philosophy of Science. Volume 7: Philosophy of Statistics.* Amsterdam: Elsevier.

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

[68] Orbán, G. & Wolpert, D.M. (2011). Representations of uncertainty in sensorimotor control. *Current Opinion in Neurobiology, 21,* 1-7.

[69] Pedersen, A.P. & Wheeler, G. (2014) Demystifying dilation. *Erkenntnis, 79,* 1305-1342.

[70] Pettigrew, R. (2016). Accuracy, risk, and the principle of indifference. *Philosophy and Phenomenological Research, 92,* 35-59.

[71] Poirier, D.J. (2006). The growth of Bayesian methods in statistics and economics since 1970. *Bayesian Analysis, 1,* 969-980.

[72] Pothos, E. M. & Busemeyer, J. R. (2013). Can quantum probability provide a new direction for cognitive modeling? *Behavioral and Brain Sciences, 36,* 255-327.

[73] Pouget, A., Beck, J.M., Ma W.J., & Latham, P.E. (2013). Probabilistic brains: knowns and unknowns. *Nature Neuroscience, 16,* 1170-1178.

[74] Rao, R.P.N., Olshausen, B.A., & Lewicki, M.S. (2002). *Probabilistic Models of the Brain: Perception and Neural Function.* Cambridge: MIT Press.

[75] Rédei, M. & Summers, S.J. (2007). Quantum probability theory. *Studies in the History and Philosophy of Modern Physics, 38,* 390-417.

[76] Rescorla, M. (2015). Bayesian perceptual psychology. in Mohan Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception.* Oxford: Oxford University Press.

[77] Rescorla, M. (2016). Bayesian sensorimotor psychology. *Mind and Language, 31,* 3-36.

[78] Rinard, S. (2015). A decision theory for imprecise probabilities. *Philosophers' Imprint, 15,* 1-16.

[79] Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: alternative algorithms for category learning. *Psychological Review, 117,* 1144.

[80] Shafer, G. (1992). The Dempster-Shafer theory. *Encyclopedia of Artificial Intelligence, 2nd edition,* New York: Wiley, 330-331.

[81] Shafer, G. & Tversky, A. (1985). Languages and designs for probability judgment. *Cognitive Science, 9,* 309-339.

[82] Simoncelli, E.P. (2009). Optimal estimation in sensory systems. in M.S. Gazzaniga (ed.), *The Cognitive Neurosciences, IV.* Cambridge: MIT Press, 525-535.

[83] Skovgaard-Olsen, N. (2016). Ranking theory and conditional reasoning. *Cognitive Science, 40,* 848-880.

[84] Sober, E. (2003). Two uses of unification. in F. Stadler (ed), *The Vienna Circle and Logical Empiricism – Vienna Circle Institute Yearbook 2002.* Dordrecht: Kluwer, 205-216.

[85] Spohn, W. (2012). *The Laws of Belief: Ranking Theory and its Philosophical Applications*. Oxford: Oxford University Press.

[86] Stanford, P.K. (2015). Unconceived alternatives and conservatism in science: the impact of professionalization, peer-review, and big science. *Synthese,* 1-18.

[87] Vilares, I. & Kording, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Annals of the New York Academy of Sciences, 1224,* 22-39.

[88] Tenenbaum, J.B., Kemp, C., Griffiths, T.L., & Goodman, N.D. (2011). How to grow a mind: statistics, structure and abstraction. *Science, 331,* 1279-1285.

[89] Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities.* London: Chapman & Hall.

[90] Williams, P. (2007). Notes on conditional previsions. *International Journal of Approximate Reasoning, 44,* 366-383.

[91] Woodward, J. (2014). Scientific explanation. *The Stanford Encyclopedia of Philosophy,* Edward N. Zalta (ed.), <http://plato.stanford.edu/archives/win2014/entries/scientific-explanation/>.

[92] Zadeh, L. (1975). Fuzzy logic and approximate reasoning. *Synthese, 30,* 407-428.

[93] Zednik, C. & Jäkel, F. (2014). How does Bayesian reverse-engineering work? in P. Bello, M. Guarini, M. McShane, & B. Scassellati (eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society.* Austin: Cognitive Science Society, 666-671.

[94] Zhao, J. & Osherson, D. (2010). Updating beliefs in light of uncertain evidence: descriptive assessment of Jeffrey's rule. *Thinking & Reasoning, 16,* 288-307.