

BAYESIAN EPISTEMOLOGY

Stephan Hartmann and Jan Sprenger

August 17, 2010

1 Introduction

Bayesian epistemology addresses epistemological problems with the help of the mathematical theory of probability. It turns out that the probability calculus is especially suited to represent degrees of belief (credences) and to deal with questions of belief change, confirmation, evidence, justification, and coherence. Compared to the informal discussions in traditional epistemology, Bayesian epistemology allows for a more precise and fine-grained analysis which takes the gradual aspects of these central epistemological notions into account. Bayesian epistemology therefore complements traditional epistemology; it does not replace it or aim at replacing it.

Bayesian epistemology can be traced back to the work of Reverend Thomas Bayes (1701-1761) who found an elementary mathematical theorem that plays a central role in Bayesian epistemology. More on this below. Later Bayesian ideas began to surface not only in philosophy, but also in statistics, formal learning theory, and other parts of science. Obviously, the probability calculus finds many applications because of its enormous flexibility, expressive power, and formal simplicity. Bayesian epistemology shares much with these endeavors, including a certain scientific attitude vis-a-vis the problems in question, but it is worth noting that Bayesian epistemology is, in the first place, a philosophical project, and that it is its ambition to further progress in philosophy.

This essay is structured as follows. Section 2 introduces the probability calculus and explains why degrees of belief obey the probability calculus. Section 3 applies the formal machinery to an analysis of the notion of evidence, and highlights potential application. Section 4 discusses Bayesian models of coherence and testimony, and section 5 ends this essay with a comparison of traditional epistemology and Bayesian epistemology.

2 Probability and Degrees of Belief

Bayesian epistemology can be described as the attempt to use an intuitive, but powerful tool – the probability calculus – for tackling long-standing problems in

epistemology and philosophy of science. In particular, Bayesian epistemology models degrees of belief as mathematical probabilities. Probability is interpreted subjectively or epistemically (as opposed to the “objective chance” of an event). This section explains the relationship between the different interpretations of probability, the concept of degree of belief, and the application of those tools to epistemological problems. We start with some motivational remarks that draw on the analogy to deductive logic.

Deductive logic is often perceived as the *logic of full rational belief*, in the sense that an agent’s set of (fully endorsed) beliefs can be described as a set of first-order propositions. If this set is logically inconsistent, i.e. if there is no joint model of all propositions, then the agent cannot be (epistemically) rational: the propositions cannot hold simultaneously, hence she ought to abandon at least one of her beliefs. The calculus of deductive logic is helpful here: it detects inconsistencies in a set of beliefs by exploring their implications according to a set of inferential rules.

The mathematical theory of probability plays the same role with respect to *partial rational beliefs*, in the sense that the probability calculus is a powerful instrument in order to infer the doxastic implications of a certain set of partial beliefs. To spell out the concept of a rational or irrational degree of belief, Frank Ramsey [1926] (1978) suggested to make use of the standard, economic conception of rationality – irrational degrees of belief would cost us money if we let them guide our actions. The crucial concept is *betting behavior*, or the inclination to accept and reject bets according to our degrees of belief. For fixed betting odds, a bet appears to be more favorable if we have a strong degree of belief in the underlying propositions than if we only weakly believe in it. This suggests to quantify a degree of belief in terms of the betting odds which we consider to be fair. As Ramsey argues, sports events are not the only occasion when we get involved in betting:

“[...] all our lives we are in sense betting. Whenever we go to the station we are betting that a train will really run, and if we had not a sufficient degree of belief in this we should decline the bet and stay at home.” (Ramsey [1926] 1978, 85.)

Thus, we can describe our degrees of belief by means of the betting odds which we consider fair. To make this explicit, we need some formalism: A *bet* on the event A is a triple $\langle A|x|y \rangle$ where x and y are positive real numbers. The bookie pays the bettor y euros if A occurs and the bettor pays the bookie x euros if A does not occur. x is the bettor’s *stake*, and the ratio $(x + y)/x$ is called the *betting odds* on A , indicating the bettor’s total gain (including the stake) for a successful €1 euro bet. Such a betting odd is a preliminary quantification of a degree of belief. Naturally, an agent judges the bet $\langle A|x|y \rangle$ to be fair if it offers no advantage to either side, i.e. if to the agent’s mind, neither the bookie

nor the bettor have an advantage, and both have the same expected utility. In this case, $(x + y)/x$ are the betting odds corresponding to the agent's degree of belief.

The connection to the probability calculus is quickly made: if x is much greater than y and the bettor takes more risk than the bookie, the agent believes the event A to be *probable*. Conversely, if agent S believes A to occur in $(100 * p)\%$ of all possible cases, he will consider the bet $\langle A|x|y \rangle$ to be fair if and only if there is no advantage for either side and the bet is a zero-sum game:

$$py + (1 - p)(-x) = 0. \tag{1}$$

The only real number that solves equation (1) is $p = x/(x + y) \in [0, 1]$. This value p is called the *probability corresponding to S 's degree of belief*. It is easy to see that there is an isomorphism between probabilities and betting odds since we can determine the probability from the betting odds by taking the inverse, and vice versa. In the remainder, we will therefore read “the probability of A ” as the subjective degree of belief in A . Evidently, 1 denotes maximal and 0 minimal degree of belief.

What exactly makes a judgment on the fairness of a bet, or actual betting behavior, irrational? If we accept 1:1 bets on an event A which you know to be highly improbable, we might still be rational. Maybe we have less information than you. Or there is not enough information available to determine a uniquely rational degree of belief. For instance, if we have a certain degree of belief in the independent propositions A and B , our degree of belief in $A \vee B$ should not be lower. Thus, degrees of belief/betting odds/probabilities are not arbitrary if we proceed from A and B to their truth-functional compounds. Probability theory reflects these entanglements in two simple axioms (Kolmogorov 1933):

Definition: Let \mathcal{A} be a field of propositions (i.e. a set of propositions that is closed under truth-functional combination and contains all tautologies). $P : \mathcal{A} \rightarrow [0, 1]$ is a *probability function* on \mathcal{A} if and only if

- $P(A) = 1$ for any tautology A .
- For incompatible (mutually exclusive) propositions A and B , $P(A \vee B) = P(A) + P(B)$.

Any such function P is called a probability. The axioms are natural: Each tautology is assigned maximal degree of belief, and the disjunction of mutually exclusive propositions is assigned the sum of the individual degrees of beliefs. As a corollary, we obtain that $P(\neg A) + P(A) = P(A \vee \neg A) = 1$.

We will see in a minute that these simple equations contain everything that an agent's rational degrees of belief have to satisfy, and vice versa. Here, the

famous *Dutch Book Theorem* comes in: if one of the probability axioms is violated, the betting odds implied by the agent's probabilities cannot have been fair altogether – it is possible to construct a system of bets that assures a risk-free gain to the bookie or the bettor, a so-called Dutch Book. Therefore these degrees of belief cannot have been rational either.

Dutch Book Theorem: Any function $P : \mathcal{A} \rightarrow [0, 1]$ on a field \mathcal{A} that does not satisfy the axioms of probability allows for a system of bets that is vulnerable to a Dutch Book.

There is a second theorem, the Converse Dutch Book Theorem, which ensures that probability functions are *not* vulnerable to Dutch Books:

Converse Dutch Book Theorem: No probability function $P : \mathcal{A} \rightarrow [0, 1]$ is vulnerable to a Dutch Book.

Proofs: See Kemeny 1955, Skyrms 1980.

In other words, the Dutch Book Theorem establishes the probability calculus as a logic of partial belief. Probabilistic degrees of belief are immune to Dutch Books, and non-probabilistic degrees of belief aren't. For instance, if your degrees of belief in A , B and $A \vee B$ did not conform to the probability axioms, we could offer you a bet on $A \vee B$ in a twofold way: directly, or as a bet implied by two single bets on A and B . But the implicit betting odds would be different, leading to a *reductio ad absurdum*:

“If anyone's mental condition violated these laws [of the probability calculus], his choice would depend on the precise form in which the options were offered him, which would be absurd.” (Ramsey [1926] 1978, 84.)

Thus, the Dutch Book Theorem establishes probability as the mathematical model of degrees of belief. However, the axiom of probabilities merely constrain *systems* of degrees of belief. They do not capture the irrationality of *isolated* beliefs. There are some people who *know* that the (objective, ontic) chance that their football club wins the national championship is no more than 0.01, but they continue to accept 1:8 bets on that event. Whatever the reasons for such a behavior, they act against their own knowledge. The inconsistency arises from the gap between the subjective degree of belief and the objective chance of the event. For this reason, we need a principle that supplements the Dutch Book Theorem with an account of the relationship between chances – objective probabilities in the world – and rational degrees of belief. This task is fulfilled by the

Principal Principle (PP, Lewis 1980): If an agent knows the objective probability of proposition A to be equal to p and he has no “overruling information” available, then his rational degree of belief in A must also be equal to p .

Lewis (1980) argues for the self-evidencing character of (PP), and this is in line with our intuitions: if we know a roulette table to be perfectly fair, we have no reason to play a specific strategy. But as Strevens (1999) points out, it is hard to give a non-circular defense of (PP).

Nevertheless, (PP) gives, together with the Dutch Book Theorem, a comprehensive account of the *statics* of rational belief. But what about the dynamics? We are lacking a principle that asserts how degrees of belief should be *changed* in the light of incoming information. Here, the third and final cornerstone of Bayesian epistemology enters, namely the updating of degrees in the light of new evidence. The degree of belief in proposition A after learning another proposition E is expressed by the conditional probability of A given E , $P(A|E)$. This value is defined as $P(A.E)/P(E)$ and is usually interpreted as the probability of A if we take E for granted.

Bayesian Conditionalization: The rational degree of belief in a proposition A after learning E is the conditional probability of A given E : $P_{\text{new}}(A) = P(A|E)$.

By means of the famous Bayes's Theorem (see Joyce 2008), we can reformulate this equation and make it easier to handle in practice:

$$\begin{aligned} P_{\text{new}}(H) = P(H|E) &= \frac{P(H)P(E|H)}{P(E)} & (2) \\ &= \frac{P(H)P(E|H)}{P(H)P(E|H) + P(\neg H)P(E|\neg H)} \end{aligned}$$

To give an easy example: A friend of yours has bought a new car. Your prior degree of belief that it is a Ford is about 0.1 (corresponding to the percentages of Fords among newly bought cars). One day, he comes to your place, driving a Ford. If he actually owns a Ford, it is quite likely that you see him in a Ford rather than in his wife's Toyota ($P(E|H) \approx 0.8$), whereas, if he did not own a Ford, he would probably taken his wife's car, public transport rather than borrowing a Ford from someone else ($P(E|\neg H) \approx 0.05$). Using (2), this leads to your new degree of belief $P_{\text{new}}(H) \approx 0.94$. In other words, you are now quite convinced that he has bought a Ford.

Together, the Dutch Book Argument, the Principal Principle and Bayesian Conditionalizations are the three pillars of Bayesian epistemology. Let us have a look what one can do with them.

3 Measuring Evidence

A central concept in modern epistemology and modern science is *evidence*. Something is evidence for a proposition or a scientific theory A , something makes us believe that A , etc. Philosophy of science has, over the past decades,

exploited the probabilistic machinery to explicate what it means that a scientific theory is confirmed or undermined. Although this debate mainly took place in philosophy of science journals, it is obviously significant for epistemology: probabilistic confirmation renders a theory more credible, in other words, it is evidence for the theory, and evidence is, in turn, central to justifications and reasons.

There are two concepts of evidence that have to be kept apart: the absolute and the relative one. According to the absolute concept, E is evidence for a proposition A if and only if, given E , A is highly probable ($P(A|E)$ is high). For instance, a perception is (absolute) evidence for a certain belief if, taking the perception or granted, the belief is highly probable. This understanding certainly captures some ways of using the word “evidence”, but on the other hand, E can be absolute evidence for A even if E lowers the probability of E . For instance, let A be the proposition that your favorite football club will *not* be national champion in the next year. Even if they win a league match (E), the probability of A given E is still sufficiently high to make E absolute evidence for A ($P(A|E)$ is still sufficiently high, although lower than $P(A|\neg E)$). At least unless you support Chelsea, Barcelona, or the like.

This unintuitive property of absolute evidence calls for a second, different concept of evidence, namely evidence in the sense of *support*. In the above example, E is evidence for $\neg A$ because it increases the chance of winning the overall competition (though only to a tiny degree). This relative concept of evidence as support is the subject of the rest of the section. Not only is it much more of a *relation* than the absolute concept, it is also fruitful and widely applicable, as we will see later.

Definition: E is (relative) evidence for a proposition A if and only if $P(A|E) > P(A)$.

Often, we have to tell good from bad evidence, similar to telling good from bad reasons, or to quantify degree of support, e.g. in order to address famous challenges such as the Duhem-Quine problem (Earman 1992). For these tasks, we need a *measure* of evidence. But what is the most adequate measure of evidence? Which one should we use when judging, for instance, whether “evidence of evidence” is also evidence? Here are some suggestions with both an intuitive appeal and a longstanding tradition (other suggestions have been made by e.g. Nozick 1981, Crupi et al. 2007):

Difference Measure Takes the difference between the posterior and the prior degree of belief in A as a measure of the support evidence E lends to A :
 $d(A, E) := P(A|E) - P(A)$. (Earman 1992, Rosenkrantz 1994)

Log-ratio Measure Proposed by Howson and Urbach (1993) and Milne (1996), this measure is based on the very same quantities, but replaces

the difference in d by the logarithmic ratio: $r(A, E) := \log P(A|E)/P(A)$.

Counterfactual Difference Measure Takes quite a different approach and compares the posterior degree of belief in A with the counterfactual degree of belief in A had $\neg E$ occurred instead of E : $s(A, E) := P(A|E) - P(A|\neg E)$. Most prominently backed by Christensen (1999) and Joyce (1999).

Log-Likelihood Ratio Measures Looks whether A or its negation $\neg A$ better accounts for observed evidence E : $l(A, E) := \log \frac{P(E|A)}{P(E|\neg A)}$.

The task of explicating evidence thus amounts to making a decision between these measures. This is an intricate task since they also capture different aspects of evidence. Let's come back to the Ford example. If you already knew that your friend had the firm plan to buy a Ford ($P(A) = 0.95$) then the evidence that you see him in a Ford is not very impressive from d 's standpoint ($d < 0.1$) since your degrees of belief do not change a lot. Still, A accounts much better for E than $\neg A$ ($l \approx 2.77$). Thus, E remains highly useful to discriminate between A and $\neg A$, and therefore strong evidence from l 's standpoint. Christensen (1999, 438-39) presents a nice analogy: How would we measure the extent to which a candidate for US presidency P is financially supported by a group G ? The proportion of G -donations in P 's funds? P 's relative position in the presidential run as a function of the G -donations? And so on. Christensen conjectures

“Thinking about these different measures of support suggests to me that there is no single clearcut question being asked when we ask ‘How much support does P get from G ?’. It would not be surprising if the same were true of the question ‘How much does evidence E support hypothesis A ?.’” (Christensen 1999, 439.)

Hence, a purely example/intuition-based approach to finding measures of evidence is misguided, since different questions are asked. The task of analyzing evidence in probabilistic terms rather consists in finding a common ground for all (relative) measures of evidence, i.e. generally accepted criteria which all measures have to satisfy. Based on such criteria, we may rule out some proposed measures and maintain others. Thus, our inability to decide specific examples, or to find conclusive counterexamples, does not entail that all measures are equally adequate. Now, we check whether certain problems are *measure-sensitive*, i.e. whether they give us a reason to choose between the various explications of evidence (Fitelson 1999, 2001a, 2001b; Eells and Fitelson 2000). The results of that discussion extend well beyond clarifying the notion of evidence: they can be used to model coherence (see the following section), to tackle longstanding philosophy of science puzzles such as the Duhem-Quine problem (see Earman 1992), and they can be connected to empirical evidence in psychology, delivering insights into human reasoning.

But let us get back to adequacy constraints on evidence measures. First, each evidence measure should, *ceteris paribus*, adequately render the force of *striking or surprising evidence*. When two pieces of evidence are equally probable given a certain hypothesis, the less expected piece of evidence seems to be more salient since we did not expect to find it. Indeed, all measures satisfy this constraint.

Fact: Assume that $P(E|A) = P(E'|A)$, $P(E) < P(E')$, and that both E and E' are evidence for A . For all evidence measures $c \in \{d,r,l,s\}$, E is better evidence for A than E' .

Thus, the problem of striking evidence is not measure-sensitive, and we have to look for other criteria in order to make a choice between the measures. What about *independent evidence* and independent testimonies? When I am standing at the bus stop and I am not sure that my watch displays the right time, I might ask someone else to have a look at her watch. I ask her because I believe the two watches to be independent of each other, and I would not ask her if I knew that we synchronized watches one hour ago. Independent evidence is not only important in such everyday situations, but also when testimonies have to be evaluated (e.g. in criminal trials), or in science, when independent replications of an experiment are supposed to confirm preliminary findings. Thus, we obtain to the following desideratum:

Desideratum (D): If E_1 and E_2 are evidence for A and if A screens off E_1 from E_2 (i.e. E_1 and E_2 are probabilistically independent of each other, conditional on A and $\neg A$, see also Reichenbach 1956) then, for an adequate measure of evidence c ,

$$c(A, E_1.E_2) > c(A, E_1). \quad (3)$$

However, Fitelson (2001a, 2001b) showed that s violates (D) – contrary to d , r and l – although (D) is arguably weak and merely requires that additional independent evidence adds *some* value to a single piece of evidence. Thus, we have a knock-down argument against s which fails to capture the force of independent evidence and testimony.

Finally, we would like to present an argument against r that has some further implications, namely the *problem of irrelevant conjunctions*. If your watch yields strong evidence that your philosophy seminar is about to begin, this should not be strong evidence for the claim that your philosophy seminar is about to begin *and* that your favorite football club is going to win the national championship. The latter proposition just seems to be irrelevant to the evidential reasoning. Formally, if E is (strong) evidence for A , then E should not always be (strong) evidence for A plus an arbitrary claim X . However, all measures of evidence yield that E is evidence for $A.X$, too, due to the positive impact of E on A :

Proposition 1 (e.g. Fitelson 2002, Hawthorne and Fitelson 2004): Assume that E is evidence for A and that $P(E|A.X) = P(E|A)$ for a sentence X consistent with A . Then E is evidence for $A.X$, too, for any measure of evidential relevance.

In other words, if we tack a conjunct to our belief in question that does not change the likelihood of the evidence, then the evidence relation extends to the conjunction as well. For instance, the seminar schedule and the precision of your watch are apparently independent of the results of football matches. And precisely such conjuncts are the ones we would call “arbitrary” or “irrelevant”, leading to the counterintuitive Proposition 1, *regardless of the used measure*. Therefore it is important that in such situations, an evidence measure \mathfrak{c} satisfies $\mathfrak{c}(A, E) > \mathfrak{c}(A.X, E)$, i.e. E is weaker evidence for the conjunction $A.X$ than for the original proposition A . It turns out that such a result is available, and that the paradox can be mitigated:

Proposition 2 (Hawthorne and Fitelson 2004): Assume that E is evidence for A and $P(E|A.X) = P(E|A)$ for a proposition X with $P(X|A.K) \neq 1$. Then the evidence E lends to A exceeds the evidence E lends to $A.X$, for d , l , and s , e.g. $d(A, E) > d(A.X, E)$. But for the log-ratio measure r , we always get $r(A, E) = r(A.X, E)$.

Thus, the measure r drops from our list of candidate measures: If E is evidence for A , then E is equally strong evidence for $A.X$ for an irrelevant X , and this is apparently an unacceptable result. Notably, the remaining measures d and l are quite opposite to each other, in the sense that d captures the (“internalist”) increase in credibility from $P(A)$ to $P(A|E)$ whereas l captures the (“externalist”) discriminative power of the evidence with respect to the propositions in question. This crucial distinction is rarely, if at all, mentioned in discussions about evidence that take place in traditional epistemology.

These results do not only affect the debate about evidence measures – they are important for the psychology of human reasoning as well. You might have heard about the *conjunction fallacy* observed by Tversky and Kahneman (1983): Take the propositions

- E : 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.
- A : Linda is a bank teller.
- $A.X$: Linda is a feminist and a bank teller.

Subjects were asked to assess whether A or $A.X$ is more probable given evidence E . Strikingly, they predominantly judged $A.X$ to be more probable, thereby

violating the probabilistic law that for each A and X , $P(A.X) \leq P(A)$. Or in other words, a proposition cannot be more credible than one of its consequences. How shall we interpret this result? Are the subjects unaware of elementary logical laws? In terms of evidential relevance and our measures of evidence, we can tell a story about their judgments (see Crupi et al. 2008, Schupbach 2009). The subjects intuitively judged the *evidential relevance* relations between E , A and $A.X$. Obviously, E is strong evidence for X (“Linda is a feminist”) and by proposition 1, at least weak evidence for $A.X$, whereas E is equally obviously evidence *against* A . This implies a fortiori that E is much stronger evidence for $A.X$ than for A . Thus, the subjects apparently confounded the concepts of high probability and the concept of evidential relevance, instead of committing a mathematical fallacy (but see the alternative story of Hartmann and Meijs 2010). This formal analysis of the problem is measure-sensitive: as proposition 2 shows, it would not be available if we used r as the measure of evidential relevance.

Similarly, claims of social epistemology can be tackled by means of evidence measures. For instance, Douven (2010) investigates whether “evidence of evidence” is evidence for a proposition, using the different explications of evidence presented in this section. All this illustrates that Bayesian techniques are not (only) a mathematician’s delight, but valuable means of tackling traditional epistemological problems. The next section describes Bayesian models of coherence and testimony.

4 Coherence and Testimony

In epistemology, the coherence theory of justification is the main alternative to foundationalism. It says that a set of propositions is justified if it coheres well (BonJour 1985). The theory is attractive as it avoids the problems of foundationalism, but it has its problems as well. Most importantly, it is not clear what it means that a set of propositions coheres. How can this notion be made more precise? The situation is complicated by the observation that coherence is a gradual notion. Some sets of propositions seem more coherent, while others are less coherent. Apparently, we need a measure that specifies how coherent a set of propositions is. Constructing such a measure may also help to get a better grasp of what coherence is. Moreover, the availability of a coherence measure will help addressing long standing problems in the coherence theory of justification. For example, one may ask if and when coherence is truth-conducive, i.e. under which conditions is the coherence of a set of propositions an indicator of its truth? Bayesian epistemologists have addressed these questions and made substantial progress over the last couple of years.

A natural way to start the construction of a measure of coherence is to

depart with an epistemological intuition and to formalize it. The following two intuitions about coherence can be identified:

(R) Coherence as positive relevance.

(O) Coherence as relative overlap in probability space.

(R) expresses the intuition that the elements of a coherent set mutually support each other (in the sense of the definition of relative evidence). Such sets of propositions seem to be more coherent than sets of independent propositions or sets whose elements are negatively relevant to each other. (O) expresses the intuition that identical propositions are considered to be coherent. This is especially plausible if one adopts a witness scenario: Imagine that several independent witnesses of a crime give identical reports (“The butler left the crime scene with a bloody knife in his hand.”). Obviously, the given reports are maximally coherent and any deviation reduces the coherence accordingly. Probabilistically speaking, identical reports maximally overlap in probability space. So coherence measures, according to (O), the relative overlap of the propositions in probability space.

Coherence measures can be classified according to which intuition they formalize. For instance, the *Shogenji measure*, the first explicit coherence measure in the literature, is a pure relevance measure (Shogenji 1999). For two propositions A and B , it is given by the following expression:

$$C_S(A, B) := \frac{P(A|B)}{P(A)} = \frac{P(B|A)}{P(B)} = \frac{P(A.B)}{P(A)P(B)} \quad (4)$$

$C_S(A, B)$ measures how relevant B is for A , i.e. how much the probability of A is raised if one learns that B is true. Note that the expression is symmetrical in A and B , which is a natural request for a coherence measure. Hence, $C_S(A, B)$ also expresses how relevant A is for B . According to this Bayesian explication, the coherence of a set of proposition is a property of this set *relative* to a probability measure P . Different agents with different probability functions may therefore come to different coherence assignments. They may also rank different sets differently according to their coherence.

The expression on the right-hand side indicates how the Shogenji measure can be generalized to more than two propositions. This generalization, however, is problematic as Fitelson (2003) has shown.

The *Glass-Olsson measure* (Glass 2002, Olsson 2005) is a pure overlap measure:

$$C_O(A, B) := \frac{P(A.B)}{P(A \vee B)} \quad (5)$$

Also $C_O(A, B)$ can be generalized to more than two propositions in a natural way.

It turns out that none of these and related measures always leads to an intuitively satisfactory coherence ordering of sets of propositions (Bovens and Hartmann 2003a, Douven and Meijs 2007, Meijs 2005, Siebel 2005). This suggests the search for more complex measures that take both intuitions – positive relevance and overlap – into account. This is achieved by the Bovens-Hartmann measure (Bovens and Hartmann 2003a) and the family of measures that generalize the Bovens-Hartmann measure (Douven and Meijs 2007). Contrary to the Shogenji measure and the Glass-Olsson measure, the construction of these measures does not start with the formalization of an epistemic intuition. It rather starts with the question what the *function* of the coherence of a set of propositions is. One obvious answer is that the coherence of a set of propositions *boosts our confidence in the truth of these propositions*. To make this approach explicit, we need to introduce a *witness scenario*.

Consider a set of propositions $S^{(n)} := \{A_1, \dots, A_n\}$. Assume that each proposition $A_i (i = 1, \dots, n)$ is confirmed by a report E_i of a different witnesses. The n witnesses are independent and have the same reliability. The construction of a coherence measure then proceeds in three steps:

1. Work out the ratio of the posterior probability $P(A_1, \dots, A_n | E_1, \dots, E_n)$ and the prior probability $P(A_1, \dots, A_n)$. The posterior probability measures the probability of the set of propositions *after* the reports came in. The prior probability measures the probability of the set of propositions *before* the reports came in. The ratio of both, then, measures the confidence boost in $S^{(n)}$. This takes intuition (R) into account. – Notably, Douven and Meijs (2007) replace the ratio measure by alternative evidence measures (e.g. d and l from section 3), thus obtaining a family of coherence measures.
2. Normalize this ratio to make sure that a set of propositions which fully overlap in probability space, has maximal coherence. This takes intuition (O) into account. It is easy to see that the resulting function cannot be a coherence measure as it depends on the reliability of the witnesses. Coherence, however, is traditionally conceived as an intrinsic property of a set of propositions and independent of the reliability of the witnesses.
3. To solve this problem, it is requested that a set $S^{(n)}$ is more coherent than a set $S'^{(n)}$ if and only if the normalized ratio of posterior and prior probability is greater for $S^{(n)}$ than for $S'^{(n)}$ *for all values* of the reliability of the witnesses. It is easy to see that this entails that there are sets of propositions $X^{(n)}$ and $Y^{(n)}$ that cannot be ordered according to their coherence, which is also intuitively plausible. See Bovens and Hartmann 2003a and 2003b for details.

A measure proposed by Fitelson (2003) also takes both intuitions into account. However, it turns out that none of the coherence measures proposed so far is without problems (Meijs and Douven 2005; Bovens and Hartmann 2005). They have all been confronted with counter-examples, i.e. with scenarios that lead to coherence orderings that are not in line with intuitions. Sometimes, however, different authors have different intuitions and no agreement is reached as to whether a certain measure is defective or not. This suggests the following conclusions: (i) One can argue that there is no measure that reflects all our intuitions about coherence. Instead one should, similarly to the debate about evidence measures, accept a plurality of coherence measures, all reflecting different aspects of coherence. (ii) One can argue that pure armchair arguments are limited. Instead of starting with various intuitions philosophers have, one should go empirical and study the coherence judgments of real people. The results can then be confronted with the coherence measures put forward in the philosophical literature. It is also possible that the data suggest a new coherence measure. For preliminary work in this direction, see Harris and Hahn 2009. See also Oaksford and Chater 2007, and Chater and Oaksford 2008. It is hoped that a combination of empirical studies, formal modeling and conceptual analysis will help to resolve the deadlock in the current debate about coherence measures.

Formal analyses of coherence have already been used to examine the relation between coherence and truth. In a witness scenario, several questions can be asked: (i) What is the effect of the coherence of a set of propositions on the posterior probability? (ii) Under which conditions is coherence truth conducive, i.e. when does the more coherent of two sets of propositions have the higher posterior probability after independent witnesses provided confirming reports? Regarding the first question, it is immediately clear that the impact of the coherence is the greater, the more independent the witnesses are. If we hear, for example, that the witnesses of a crime scene talk extensively prior to a police interrogation, then we are less impressed by their coinciding reports. In this case, the coherence will not increase the posterior probability much. Generally speaking, coherence is maximally effective if the witnesses are independent.

Hints to answering the second question come from two impossibility theorems, which show that there is no coherence measure that fulfills a number of plausible requirements. While Olsson (2005) concludes on the basis of his impossibility theorem that coherence is no indicator for truth, Bovens and Hartmann (2003a, 2005, 2006) argue that the negative result of their impossibility theorem can be avoided if one does not require that sets of propositions can *always* be ordered according to their coherence. However, in certain cases coherence is an indicator of truth in the following sense: if two sets of propositions of equal cardinality have the same prior probability and can be ordered according to their coherence, then the more coherent of the two also has the higher posterior

probability if all witnesses have the same reliability.

There are still many open questions that will be addressed in future work. Here are two: First, more realistic models of information gathering have to be constructed. One step in this direction is to relax the independence condition. Can one then still prove an interesting impossibility theorem? Second, the use of coherence considerations in belief revision needs to be considered in more detail. These results will also be of relevance for philosophy of science (theory change) and computer science.

We conclude this section with a discussion of testimony. Epistemologists have stressed that much of our knowledge derives from the testimony of others, parents, teachers, textbooks etc. While this seems true, more specific questions can be asked. For example, how shall the testimonies of several witnesses be combined? And how shall we change our beliefs in the light of testimonial evidence? Bayesian epistemology has the resources to make these questions more precise and to answer them. In *Bayesian Epistemology*, Bovens and Hartmann (2003a) develop a general methodology, using the theory of Bayesian networks (Neapolitan 2003), that facilitate a detailed analysis.

More specifically, models with more or less dependent witnesses of different reliability can be considered and a range of interesting results can be proven. For example, “too-odd-to-be-true” reasoning can be studied: imagine that several independent and partially reliable witnesses give the same report; they all claim that they saw no. 1 in a lineup of n people at the crime scene. We ask: When are we more convinced that the witnesses tell the truth, if the number of suspects is, say, 2 or if it is 100? Most people would say that we are the more convinced in the truth of the reports, the larger the number of suspects is. This makes much sense as it becomes increasingly unlikely that a witness hits the truth by accident if the number of suspects is large. The probability for this is only $1/n$ and decreases if n goes up. From a Bayesian perspective, however, different perspectives on this phenomenon are possible: On the one hand, coinciding reports provide the better relative evidence (cf. section 3) the higher the number of suspects is. This seems to confirm our intuitive judgment. But on the other hand, with increasing number of suspects, the prior probability of no. 1 being responsible for the crime declines. Thus the posterior probability of the witnesses telling the truth falls, too, and the coincidence is “too odd to be true”. (Note the analogy to the Linda case!) This example illustrates once more that Bayesian epistemology provides fruitful tools to tackle questions of coherence, testimony and reliable evidence. For a detailed models and discussion, see Olsson 2002 and Bovens and Hartmann 2003a, ch. 5.

5 Concluding Remarks

We conclude this essay with a comparison of traditional epistemology and Bayesian epistemology and a few remarks about the future of epistemology. Traditional epistemology typically starts with an *epistemic intuition*, developed, perhaps, by examining an example in some detail. Think about the Gettier cases as an illustration. These intuitions inspire a philosophical theory which, in turn, is criticized by other examples, triggered by different or more fine-grained intuitions.

Bayesian epistemology, on the other hand, draws much of its power from the mathematical machinery of probability theory. It starts with a *mathematical intuition*. The construction of Bayesian models is much triggered by what is mathematically elegant and feasible (e.g. Spirtes et al. 2001). The mathematics develops a life of its own (to adopt a phrase due to Hacking), and the comparison with intuitive examples comes only *after* the Bayesian account is given.

One of the goals of this essay was to show that traditional epistemology and Bayesian epistemology can learn from each other: Bayesian epistemology makes certain debates in traditional epistemology more precise, and traditional epistemology inspires Bayesian accounts.

Both Bayesian epistemology and traditional epistemology do not much consider empirical data. Both are based on intuitions, be they mathematical or epistemological. This might be a problem as privilege is given to the philosopher's intuitions. However, recent work in experimental philosophy (Knobe and Nichols 2008) has shown that non-philosophers may have different intuitions. While it is debatable how serious these intuitions should be taken (maybe people are simply wrong!), it seems clear that both can profit from taking empirical studies into account so that epistemology becomes, at the end, an endeavor to which philosophers with various tools (conceptual analysis, formal methods, and empirical studies) contribute.

The Bayesian framework is very convenient for these studies. It is easy to use, very powerful, and it comes with principles that can be made plausible on rational grounds (Dutch book arguments, Principal Principle, Bayesian Conditionalization). Moreover, this framework has been very successful as the various applications in epistemology and philosophy of science demonstrate. It has been a progressive research programme, to use Lakatos's terminology. But the Bayesian framework may reach its limits just like a scientific theory has to be given up at some point. It is therefore advisable that philosophers also keep on paying attention to other formal frameworks such as alternative theories of uncertainty (such as the Dempster-Shafer theory; see Haenni and Hartmann 2006) and epistemic logics.

References

- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge/MA: Harvard University Press.
- Bovens, L. and S. Hartmann (2003a). *Bayesian Epistemology*. Oxford: Oxford University Press.
- Bovens, L. and S. Hartmann (2003b), Solving the Riddle of Coherence, *Mind* 112, 601–634.
- Bovens, L. and S. Hartmann (2005). Why There Cannot be a Single Probabilistic Measure of Coherence, *Erkenntnis* 63, 361–374.
- Bovens, L. and S. Hartmann (2006). An Impossibility Result for Coherence Rankings, *Philosophical Studies* 128, 77–91.
- Chater, N. and M. Oaksford (eds.) (2008). *The Probabilistic Mind: Prospects for Bayesian Cognitive Science*. Oxford: Oxford University Press.
- Christensen, D. (1999). Measuring Confirmation. *Journal of Philosophy* 96: 437–461.
- Crupi, V., K. Tentori and M. Gonzalez (2007). On Bayesian Measures of Evidential Support: Theoretical and Empirical Issues. *Philosophy of Science* 74(2): 229–252.
- Crupi, V., B. Fitelson and K. Tentori (2008). Probability, Confirmation and the Conjunction Fallacy. *Thinking and Reasoning* 14(2): 182–199.
- Douven, I. (2010). Is evidence of evidence evidence? *Under review*.
- Douven, I. and W. Meijs (2007). Measuring Coherence. *Synthese* 156: 405–425.
- Earman, J. (1992). *Bayes or Bust?* Cambridge/MA: The MIT Press.
- Eells, E. and B. Fitelson (2000). Measuring Confirmation and Evidence. *Journal of Philosophy* 97: 663–672.
- Fitelson, B. (1999). The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity. *Philosophy of Science* 66: S362–S378.
- Fitelson, B. (2001a). *Studies in Bayesian Confirmation Theory*. PhD thesis, University of Wisconsin, Madison.
- Fitelson, B. (2001b). A Bayesian Account of Independent Evidence with Applications. *Philosophy of Science* 68: S123–S140.

- Fitelson, B. (2002). Putting the Irrelevance Back into the Problem of Irrelevant Conjunction. *Philosophy of Science* 69: 611–622.
- Fitelson, B. (2003). A Probabilistic Theory of Coherence. *Analysis* 63: 194–199.
- Glass, D. (2002). Coherence, Explanation, and Bayesian Networks. In: M. O’Neill et al. (eds.), *AICS 2002, LNAI 2464*, 177–182.
- Haenni, R. and S. Hartmann (2006). Modeling Partially Reliable Information Sources: A General Approach Based on Dempster-Shafer Theory, *Information Fusion* 7, 361–379.
- Harris, A.J.L. and Hahn, U. (2009). Bayesian Rationality in Evaluating Multiple Testimonies: Incorporating the Role of Coherence, *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 35: 1366–1373.
- Hartmann, S. and W. Meijs (2010). Walter the Banker: The Conjunction Fallacy Reconsidered, to appear in *Synthese*.
- Harris, A.J.L. and Hahn, U. (2009). Bayesian Rationality in Evaluating Multiple Testimonies: Incorporating the Role of Coherence, to appear in: *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Howson, C. and P. Urbach (1993). *Scientific Reasoning: The Bayesian Approach*. La Salle: Open Court.
- Joyce, J. (1999). *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Joyce, J. (2008). Bayes’ Theorem. *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), E. Zalta (ed.).
- Kemeny, J. G. (1955). Fair Bets and Inductive Probabilities. *Journal of Symbolic Logic* 20: 263–273
- Knobe, J. and S. Nichols (2008). *Experimental Philosophy*. New York: Oxford University Press.
- Kolmogorov, A.N. (1933). *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Berlin: Springer. English translation: *Foundations of the Theory of Probability*. New York: Chelsea 1950.
- Lewis, D. (1980). A Subjectivist’s Guide to Objective Chance. In R. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, Vol II., Berkeley: University of California Press, 263–293.

- Meijs, W. (2005). *Probabilistic Measures of Coherence*. Ph.D. thesis, University of Rotterdam.
- Meijs, W. and I. Douven (2005). Bovens and Hartmann on Coherence. *Mind* 114: 355–363.
- Milne, P. (1996). $\log[p(h/eb)/p(h/b)]$ is the one true measure of confirmation. *Philosophy of Science* 63: 21–26.
- Neapolitan, R. (2003). *Learning Bayesian Networks*. Prentice Hall.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Oaksford, M. and N. Chater (2007). *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford: Oxford University Press.
- Olsson, E. (2002). Corroborating Testimony, Probability and Surprise, *The British Journal for the Philosophy of Science* 53(2): 273–288.
- Olsson, E. (2005). *Against Coherence: Truth, Probability, and Justification*. Oxford: Oxford University Press.
- Ramsey, F.P. (1978). Truth and Probability, in Hugh Mellor (ed.): *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*, 58–100. London: Routledge. Original article published in 1926.
- Reichenbach, H. (1956). *The Direction of Time*. Berkeley: University of California Press.
- Rosenkrantz, R. D. (1994). Bayesian Confirmation: Paradise Regained. *British Journal for the Philosophy of Science* 45: 467–476.
- Schupbach, J. (2009). Is the Conjunction Fallacy tied to Probabilistic Confirmation?, to appear in *Synthese*.
- Shogenji, T. (1999). Is Coherence Truth-Conducive? *Analysis* 59: 338–345.
- Skyrms, B. (1980). *Causal Necessity*. New Haven: Yale University Press.
- Siebel, M. (2005). Against Probabilistic Measures of Coherence. *Erkenntnis* 63 (3): 335–360.
- Spirtes, P., C. Glymour, and R. Scheines (2001). *Causation, Prediction and Search*. Cambridge, Mass.: MIT Press.
- Strevens, M. (1999). Objective Probability as a Guide to the World. *Philosophical Studies* 95: 243–275.

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

Further readings

Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: Chicago University Press. (A classic, nowadays out-of-fashion theory of evidence.)

Fitelson, B. (2001a). *Studies in Bayesian Confirmation Theory*. PhD thesis, University of Wisconsin, Madison. (Comprehensive overview of the debate on measures of evidence.)

Jeffrey, R. (1983). *The Logic of Decision*. 2nd ed. Chicago: University of Chicago Press. (Explicates Jeffrey Conditionalization – a more general version of Bayesian Conditionalization that allows to conditionalize on uncertain evidence.)

Williamson, J. (2005). *Bayesian Nets and Causality: Philosophical and Computational Foundations*. Oxford: Oxford University Press. (An introduction to objective Bayesianism – the position that only one degree of belief is rational in a given information situation.)

Notes on Contributors

Stephan Hartmann (Ph.D. (Philosophy) at the University of Gießen, 1995) is Chair of Epistemology and Philosophy of Science at Tilburg University and founding director of the Tilburg Center for Logic and Philosophy of Science (since 2007). Stephan has been publishing extensively on philosophy of science and the use of formal methods in epistemology.

Jan Sprenger (Ph.D. (Philosophy) at the University of Bonn, 2008) is Assistant Professor of Philosophy at Tilburg University (since 2008). Jan’s work focuses on philosophy of statistics, philosophy of probability, decision theory and formal modeling in social epistemology.