

Solving the Riddle of Coherence

LUC BOVENS AND STEPHAN HARTMANN

A *coherent* story is a story that fits together well. This notion plays a central role in the coherence theory of justification and has been proposed as a criterion for scientific theory choice. Many attempts have been made to give a probabilistic account of this notion. A proper account of coherence must not start from some partial intuitions, but should pay attention to the *role* that this notion is supposed to play within a particular context. Coherence is a property of an information set that boosts our confidence that its content is true *ceteris paribus* when we receive information from independent and partially reliable sources. We construct a measure c_r that relies on hypothetical sources with certain idealized characteristics. A maximally coherent information set, that is, a set with equivalent propositions, affords a maximal confidence boost. c_r is the ratio of the actual confidence boost over the confidence boost that we would have received, had the information been presented in the form of maximally coherent information, *ceteris paribus*. This measure is functionally dependent on the degree of reliability r of the sources. We use c_r to construct a coherence quasi-ordering over information sets S and S' : S is no less coherent than S' just in case $c_r(S)$ is not smaller than $c_r(S')$ for *any* value of the reliability parameter. We show that, on our account, the coherence of the story about the world gives us a reason to believe that the story is true and that the coherence of a scientific theory, construed as a set of models, is a proper criterion for theory choice.

When we are curious about what some corner of the world looks like, we can do all kinds of things to find an answer. We can go take a look ourselves, ask experts, consult textbooks ... Sometimes the story that we get falls nicely into place like a jigsaw puzzle and we are confident that the story is true. Sometimes the story that we get just does not hang together and we are doubtful that the story is true. Some philosophers have been very impressed by this item of common sense and have made it into the cornerstone of their coherence theory of justification. Consider the grand sceptical question: do we have any reason to believe the story about the world that materialized through centuries of (scientific) inquiry? Common sense teaches us that if the story that we gather about some small corner of the world is a coherent story, then this gives us at least some reason to believe it. Some philosophers propose to recruit this item of common sense to combat sceptical worries: it is the very coherence of the story about the world that warrants our confidence.

When we borrow from common sense, it is important that we understand what common sense really has to offer. Let us represent a story by an information set, that is, a set of propositions that we have acquired through gathering information. We can give examples of information sets that are more coherent than other information sets. BonJour presents us with the following two sets: $S = \{[\text{All ravens are black}]^1, [\text{This bird is a raven}], [\text{This bird is black}]\}$ and $S' = \{[\text{This chair is brown}], [\text{Electrons are negatively charged}], [\text{Today is Thursday}]\}$ (1985, p. 96). Set S is clearly more coherent than set S' . But this is no more than an example. There is a long-standing embarrassment here. A definition of what it means for one set of propositions to be more coherent than another set has not been forthcoming. Already in 1934, A. C. Ewing writes that the absence of such a definition reduces the theory 'to the mere uttering of a word, coherence, which can be interpreted so as to cover all arguments, but only by making its meaning so wide as to rob it of almost all significance' (1934, p. 246). In *The Structure of Empirical Knowledge*, BonJour says many illuminating things about coherence, but admits that his account 'is a long way from being as definitive as desirable' (1985, p. 101) and more recently he writes that 'the precise nature of coherence remains an unsolved problem' (1999, p. 124).

The notion of coherence also plays a role in the philosophy of science. Kuhn (1977, pp. 321–2) mentions *consistency* as one of the (admittedly imprecise) criteria for scientific theory choice, along with accuracy, scope, simplicity and fruitfulness. Salmon (1990, p. 198) distinguishes between the internal consistency of a theory and the consistency of a theory with other accepted theories. In discussing the latter type of consistency, he claims that there are two aspects to this notion, namely, the '*deductive* relations of entailment and compatibility' and the '*inductive* relations of fittingness and incongruity'. We propose to think of the internal consistency of a theory in the same way as Salmon thinks of the consistency of a theory with accepted theories. Hence, the *internal consistency* of a theory matches the epistemologist's notion of the *coherence* of an information set: how well do the various components of the theory fit together, how congruous are these components? Salmon also writes that this criterion of consistency 'seem[s] ... to pertain to assessments of the prior probabilities of the theories' and 'cr[ies] out for a Bayesian interpretation' (1990, p. 198).

¹ Following Quine (1960, p. 168), the square brackets are used to refer to the proposition expressed by the enclosed sentence.

Following Salmon's lead, we will construct a probabilistic measure that permits us to read off a coherence quasi-ordering over a set of information sets and show how our relative degree of confidence that one or another scientific theory is true is functionally dependent on this quasi-ordering.²

1. Strategy

The problem with existing probabilistic accounts of coherence is that they try to bring precision to our intuitive notion of coherence independently of the particular role that it is meant to play within the coherence theory of justification.³ This is a mistake. To see this, consider the following analogy. We not only use the notion of coherence when we talk about information sets, but also, for example, when we talk about groups of individuals. Group coherence tends to be a good thing. It makes ant heaps more fit for survival, it makes law firms more efficient, it makes for happier families, and so on. There is not much sense in asking what makes for a more coherent group independently of the particular role that coherence is supposed to play in the context in question. We must first fix the context in which coherence purports to play a particular role. For instance, let the context be ant heaps and let the role be that of promoting reproductive fitness. We give more precise content to the notion of coherence in this context by letting coherence be the property of ant colonies that plays the role of boosting fitness and at the same time matches our pre-theoretic notion of the coherence of social units. A precise fill-in for the notion of coherence will differ as we consider fitness boosts for ant heaps, efficiency boosts for law firms, or happiness boosts for families.

Similarly, it does not make any sense to ask what precisely makes for a more coherent information set independently of the particular role that coherence is supposed to play within the context in question. What is this context and what is this role? The coherence theory of justification rides on a particular common sense intuition. When we gather information, then the more coherent the story that materializes is, the more confident we may be, *ceteris paribus*. In other words, within the context of information gathering, coherence is a property of information sets that plays a confidence boosting role. But what features should

² Following Sen (1970, p. 9), a quasi-ordering (or a pre-ordering) is a binary relation that is transitive and reflexive, but need not be complete, which distinguishes it from an ordering.

³ Lewis (1946, p. 338) is the *locus classicus*. More recently, proposals were suggested or defended by Shogenji (1999), Olsson (2002, p. 250) and Fitelson (2003).

the information sources have, so that the coherence of the information set indeed boosts our degree of confidence that the information is true? And what goes into the *ceteris paribus* clause—what other factors affect our confidence that the information is true? Let us take up these questions in turn.

2. The nature of the information sources⁴

Coherence will play a confidence boosting role when the information sources are independent and partially reliable. There are two aspects to independence. C. I. Lewis writes that ‘it is the possible role of congruence [that is, coherence] in the determination of empirical truth which is dramatized in detective stories ...’ (1946, p. 346). So let us construct a detective story. We interview various sources in a murder case and the information set that materializes fits nicely together: one source tells us that the murderer was wearing a butler’s jacket, a second source that the murder was committed with a kitchen knife and a third source reports that the butler was having an affair with the victim’s wife. First, we will not be impressed by the coherence of this story if the sources have had a chance to put their heads together and concoct a coherent story. This would make us feel little more confident than if we had read the story in the yellow press. Second, we will not be impressed by the coherence of this story if it turns out that each source has witnessed precisely the same fact, namely, that the murderer was wearing a butler’s jacket. The first source simply reported this fact, but the second and third sources respectively inferred that the murder weapon was probably a kitchen knife and that the butler was probably having an affair with the victim’s wife, and they chose to report this instead. The coherence of the story is of no consequence when the sources have had a chance to confer or when the sources are reporting what they inferred from the facts that other sources are reporting on. So a plausible interpretation of independence in this context is that independent sources are sources that gather information by and only by observing the facts that they report on. They may not always provide a correct assessment of these facts, but they are not influenced by the reports of other sources, nor by the facts that other sources are reporting on.

⁴ Our model of independent sources can be found in Bovens and Olsson (2000, pp. 690 and 696–70 and 2002, pp. 143–4) and in Earman (2000, pp. 56–9). Our model of partially reliable sources matches interpretation (ii) of ‘dubious information gathering processes’ in Bovens and Olsson (2000, p. 698). Olsson (2002, p. 259) substitutes ‘partial reliability’ for ‘relative unreliability’—which is the term used by Lewis (1946, p. 346).

What can we say about the reliability of the sources? First, if we know that the information sources are fully reliable, that is, that they are truth-tellers, then the coherence of the information set is a red herring. Whether the story that materializes hangs together nicely or not, the content of the story is certain. Second, if we know that the information sources are randomizers, that is, that they just flip a coin to determine how they will answer your questions, then the information that we receive is entirely useless. If some coherent story were to materialize, then it would be entirely due to chance. The coherence of the story does not increase our confidence in the story a bit. So we take our sources to be partially reliable—they are more reliable than randomizers, but short of being fully reliable.

3. The determinants of our degree of confidence

Suppose that we receive items of information from independent and partially reliable sources, say, observations, witness consultations, experimental tests, and so on. Then what determines our degree of confidence that the conjunction of these items of information is true? Consider the following case. We are trying to determine the locus of the faulty gene on the human genome that is responsible for a particular disease. Before conducting the experiments, there are certain loci that we consider to be more likely candidates. We run tests with independent and partially reliable instruments. Each test identifies an area on the human genome where the faulty gene might be located. It turns out that there is a certain overlap between the indicated areas. It is plausible that the following three factors affect our degrees of confidence that our tests are providing us with correct data, that is, that the faulty gene is indeed located somewhere in the overlapping area.

(i) How expected are the results?

Compare two cases of the above procedure. Suppose that the only difference between the cases is that, given our background knowledge, in one case the overlapping area is initially considered to be a highly expected candidate area, whereas in the other case the overlapping area is initially considered to be a highly unexpected candidate area for the faulty gene. Then clearly, our degree of confidence that the locus of the faulty gene is in the overlapping area will be lower in the latter than in the former case.

(ii) How reliable are the tests?

Again, compare two cases of the above procedure. Suppose that the only difference between the cases is that in one case the tests are highly reliable, whereas in the other case they are highly unreliable. Then clearly, our degree of confidence that the locus of the faulty gene is in the overlapping area will be lower in the latter than in the former case.

(iii) How coherent is the information?

This time suppose that the only difference is that in one case the test reports fit together extremely well—they point to precisely the same relatively narrow area—whereas in the other case they fit together poorly—they point to broad and diverging areas but the overlap between these areas coincides with the relatively narrow area in the first case. Then clearly, our degree of confidence that the locus of the faulty gene is in the overlapping area will be lower in the latter than in the former case.

We will construct a model in order to define a measure for each of these determinants. It is easy to construct an expectance measure and a reliability measure. As to coherence, the matter is more complex. We will show that there does not exist a coherence measure as such, but will define a measure that yields a coherence quasi-ordering over information sets.

4. The model

Suppose that there are n independent and relatively unreliable sources and source i informs us of a proposition R_i , for $i = 1, \dots, n$, so that the information set is $\{R_1, \dots, R_n\}$. For proposition R_i (in roman script) in the information set, let us define a propositional variable R_i (in italic script) which can take on two values, namely, R_i and $\neg R_i$, for $i = 1, \dots, n$. Let $REPR_i$ be a propositional variable which can take on two values, namely, $REPR_i$, that is, upon consultation with the proper source, there is a report to the effect that R_i is the case, and $\neg REPR_i$, that is, upon consultation with the proper source, there is no report to the effect that R_i is the case. We construct a joint probability distribution P over $R_1, \dots, R_n, REPR_1, \dots, REPR_n$, satisfying the constraint that the sources are independent and partially reliable.

We model our earlier account of the independence of the sources by stipulating that P respects the following conditional independences:

- (1) $REPR_i \perp\!\!\!\perp R_1, REPR_1, \dots, R_{i-1}, REPR_{i-1}, R_{i+1}, REPR_{i+1}, \dots, R_n, REPR_n | R_i$ for $i = 1, \dots, n$ (see footnote⁵)

or, in words, $REPR_i$ is probabilistically independent of $R_1, REPR_1, \dots, R_{i-1}, REPR_{i-1}, R_{i+1}, REPR_{i+1}, \dots, R_n, REPR_n$, given R_i , for $i = 1, \dots, n$. What this means is that the probability that we will receive a report that R_i given that R_i is the case (or is not the case) is not affected by any additional information about whether other propositions are true or whether there is a report to the effect that other propositions are true. Equivalently, in the terminology of the theory of probabilistic causation, we say that R_i screens off $REPR_i$ from all other fact variables R_j and from all other report variables $REPR_j$.⁶

We make the assumption that our partially reliable sources are all equally reliable. We specify the following two parameters, that is to say, the true positive rate $P(REPR_i | R_i) = p$ and the false positive rate $P(REPR_i | \neg R_i) = q$ for $i = 1, \dots, n$. We will discuss the assumption of equal reliability in section 7. If the information sources are truth-tellers, then $q = 0$. They never provide a report to the effect that R_i given that R_i is false. If they are fully unreliable, then it is as if they did not even attend to whether R_i is or is not the case. Rather they act as if they were randomizers. They might as well flip a coin or cast a die to determine whether they will or will not report that R_i is the case. Hence, for fully unreliable information sources, $p = q$ for $p, q \in (0, 1)$. Since partially reliable information sources are more reliable than randomizers,⁷ but less reliable than truth-tellers, we impose the following constraint on P :

- (2) $p > q > 0$

The degree of confidence in the information set is the posterior joint probability of the propositions in the information set after all the reports have come in:

- (3) $P^*(R_1, \dots, R_n) = P(R_1, \dots, R_n | REPR_1, \dots, REPR_n)$.

⁵ This notation was introduced by Dawid (1979) and has become standard notation. See Pearl (2000) and Spirtes *et al.* (2000).

⁶ See Reichenbach (1956) and Salmon (1998).

⁷ One might object that liars are even less reliable than randomizers since they consistently provide us with false information. But notice that we assume in our model that the values of p and q are known. Our sources are like medical tests whose false positive rates q and false negative rates $1 - p$ are known. Now suppose that a source is a consistent liar, that is, $p = 0$ and $q = 1$. Such a source actually becomes a very reliable source, since if the source reports that R_i then we can be certain that R_i is false and if the source refrains from reporting that R_i then we can be certain that R_i is true.

We apply Bayes' Theorem and expand the denominator by means of the chain rule:

$$(4) \quad P^*(R_1, \dots, R_n) = \frac{P(\text{REPR}_1, \dots, \text{REPR}_n | R_1, \dots, R_n) P(R_1, \dots, R_n)}{\sum_{R_1, \dots, R_n} P(\text{REPR}_1, \dots, \text{REPR}_n | R_1, \dots, R_n) P(R_1, \dots, R_n)}$$

We simplify on grounds of the conditional independences in (1):

$$(5) \quad P^*(R_1, \dots, R_n) = \frac{P(\text{REPR}_1 | R_1) \times \dots \times P(\text{REPR}_n | R_n) \times P(R_1, \dots, R_n)}{\sum_{R_1, \dots, R_n} P(\text{REPR}_1 | R_1) \times \dots \times P(\text{REPR}_n | R_n) \times P(R_1, \dots, R_n)}$$

Let $P(R_1, \dots, R_n) = a_0 > 0$.⁸ The subscript 'o' indicates that we assess the prior probability of the conjunction of $n - 0$ positive values and 0 negative values of the variables R_1, \dots, R_n . Hence, the numerator of (5) equals:

$$(6) \quad P(\text{REPR}_1 | R_1) \times \dots \times P(\text{REPR}_n | R_n) \times P(R_1, \dots, R_n) = p^n a_0$$

In the denominator, we will collect the terms in the sum in the following way. First, we calculate the term in which the variables R_1, \dots, R_n take on n positive values and 0 negative values. This is once again the expression in (6). Subsequently we collect the terms in which the variables R_1, \dots, R_n take on $n - 1$ positive values and 1 negative value. Let us consider one such term, namely, the term in which the variable R_1 takes on a negative value and R_2, \dots, R_n take on positive values:

$$(7) \quad P(\text{REPR}_1 | \neg R_1) \times P(\text{REPR}_2 | R_2) \times \dots \times P(\text{REPR}_n | R_n) \times P(\neg R_1, \dots, R_n) = qp^{n-1} P(\neg R_1, \dots, R_n)$$

We do the same for all the other terms in which the variables R_1, \dots, R_n take on $n - 1$ positive values and 1 negative value:

$$(8) \quad P(\text{REPR}_1 | R_1) \times P(\text{REPR}_2 | \neg R_2) \times \dots \times P(\text{REPR}_n | R_n) \times P(R_1, \neg R_2, \dots, R_n) = qp^{n-1} P(R_1, \neg R_2, \dots, R_n)$$

(...)

$$(9) \quad P(\text{REPR}_1 | R_1) \times \dots \times P(\text{REPR}_n | \neg R_n) \times P(R_1, \dots, \neg R_n) = qp^{n-1} P(R_1, \dots, \neg R_n)$$

Let $P(\neg R_1, \dots, R_n) + \dots + P(R_1, \dots, \neg R_n) = a_1$. The subscript '1' indicates that we assess the sum of the prior probabilities of the conjunctions of $n - 1$

⁸ If $P(R_1, \dots, R_n) = 0$, then the information is inconsistent, and however reliable our sources may be, $P^*(R_1, \dots, R_n) = 0$, that is, our confidence in the information remains unaffected. We take inconsistency to be a limiting case of a lack of coherence.

positive values and 1 negative value of the variables R_1, \dots, R_n . Hence, a_1 is the prior probability that exactly one proposition in the information set is false. We can now construct the sum of the expressions in (7) through (9):

$$(10) \quad qp^{n-1}P(\neg R_1, \dots, R_n) + \dots + qp^{n-1}P(R_1, \dots, \neg R_n) = qp^{n-1}a_1$$

Subsequently we gather all terms in the denominator in which the variables R_1, \dots, R_n take on $n - 2$ positive values and 2 negative values and so on until we reach the term in which these variables take on 0 positive values and n negative values. We can now write down the denominator in (5):

$$(11) \quad \sum_{R_1, \dots, R_n} P(\text{REPR}_1 | R_1) \times \dots \times P(\text{REPR}_n | R_n) \times P(R_1, \dots, R_n) = p^n a_0 + qp^{n-1}a_1 + q^2p^{n-2}a_2 + \dots + q^n a_n$$

and hence,

$$(12) \quad P^*(R_1, \dots, R_n) = \frac{p^n a_0}{p^n a_0 + qp^{n-1}a_1 + q^2p^{n-2}a_2 + \dots + q^n a_n}.$$

Divide numerator and denominator by p^n in (12) and substitute x for the likelihood ratio q/p :

$$(13) \quad P^*(R_1, \dots, R_n) = \frac{a_0}{\sum_{i=0}^n a_i x^i}$$

To illustrate (13), we constructed a diagram which represents a joint probability distribution over the propositional variables R_1, R_2, R_3 and contains the corresponding values for a_i , for $i = 0, \dots, 3$ in Figure 1. Let us name $\langle a_0, \dots, a_n \rangle$ the *weight vector* of the information set $\{R_1, \dots, R_n\}$. Note that:

$$\sum_{i=0}^n a_i = 1.$$

Suppose that the sources are twice as likely to report that R_i is the case when it is the case than when it is not the case, so that $x = .5$. Then our degree of confidence after we have received the reports from the sources is:

$$(14) \quad P^*(R_1, \dots, R_n) = \frac{.05}{.05 \times .5^0 + .30 \times .5^1 + .45 \times .5^2 + .20 \times .5^3} \approx .15$$

$$\begin{aligned}
 a_0 &= .05 \\
 a_1 &= 3 \times .10 = .30 \\
 a_2 &= 3 \times .15 = .45 \\
 a_3 &= .20
 \end{aligned}$$




Figure 1 Calculating the weight vector for a probability distribution over R_1 , R_2 , and R_3

5. Expectance, reliability and coherence

We can directly identify the first determinant of the degree of confidence in the information set. Note that $a_0 = P(R_1, \dots, R_n)$ is the prior joint probability of the propositions in the information set, that is, the probability before any information was received. This prior probability is higher for more expected information and lower for less expected information. Let us call this prior probability the *expectance measure*. Note that $P^*(R_1, \dots, R_n)$ increases as we increase a_0 and decrease at least one a_i (for $i \in \{1, \dots, n\}$) so that $\sum_{i=0}^n a_i$ remains 1.

We can also directly identify the second determinant, that is, the reliability of the sources. Note that $P^*(R_1, \dots, R_n)$ in (13) is a monotonically decreasing function of the likelihood ratio $x = q/p$. Let us call $r := 1 - x$ the *reliability measure*. $P^*(R_1, \dots, R_n)$ is a monotonically increasing function of r and the limits of this measure are 0 for sources that are randomizers and 1 for sources that are truth-tellers.⁹

⁹Note that r measures the reliability of the source with respect to the report in question and not the reliability of the source *tout court*. To see this distinction, consider the case in which q equals 0. In this case, r reaches its maximal value 1, no matter what the value of p is. Certainly, a source that provides fewer rather than more false negatives, as measured by $1 - p$, is a more reliable source *tout court*. But when q is 0, the reliability of the source with respect to the report in question is not affected by the value of $p > 0$. No matter what the value of p is, we can be fully confident that what the source says is true, since $q = 0$ —that is, she never provides any false positives. We will use the elliptical expression of the reliability of the source to stand for the reliability of the source *with respect to the report in question*, not for the reliability of the source *tout court*.

Let us now turn to the third determinant, namely, the coherence of the information set. The coherence of the information set is some function of the *weight vector* $\langle a_0, \dots, a_n \rangle$ of the information set $\{R_1, \dots, R_n\}$. A maximally coherent information set has the weight vector $\langle a_0, 0, \dots, 0, 1 - a_0 \rangle$. In this case, all items of information R_1, \dots, R_n are equivalent. If a_1 or $a_2 \dots$ or a_{n-1} exceeds 0, then the items of information are no longer equivalent and the information set loses its maximal coherence. But it is not clear what function of the weight vector determines the coherence of the information set.

We mentioned earlier that coherence is the confidence-boosting property of information sets. We could measure this confidence boost by considering the ratio

$$(15) \quad b(\{R_1, \dots, R_n\}) = \frac{P^*(R_1, \dots, R_n)}{P(R_1, \dots, R_n)} .$$

But this runs into the following problem. Suppose that we have a set of propositions that are independent of each other whose prior joint probability is rather low, say, .10, and a set of propositions that are equivalent to each other whose prior joint probability is rather high, say, .90. For a particular value of the reliability parameter r , the posterior joint probability of the propositions in the former set may double, but the posterior joint probability of the propositions in the latter set can maximally increase by a factor 10%. And clearly, a set of equivalent propositions is more coherent than a set of independent propositions.

Our strategy will be to assess the coherence of an information set by measuring the proportion of the confidence boost b that we actually receive, relative to the confidence boost b^{\max} that we would have received, *had we received this very same information in the form of maximally coherent information*. So we measure the proportional confidence-boosting property of the information set. The following example will make this clear. Consider once again our example of independent tests that identify sections on the human genome that contain the locus of a genetic disease. The tests pick out different areas and the overlap between these areas is σ . The information is more coherent when the reports are all clustered around the region σ than when they are scattered all over the human genome but have this relatively small area of overlap on the region σ . The information is maximally coherent when every single test points to the region σ . Suppose that we had received the information from our tests in the form of maximally coherent information, *ceteris paribus*. We calculate the degree of confidence that

the locus of the disease is in the region σ . Let R_i^σ be the proposition that test i picks out the area σ for $i = 1, \dots, n$. We construct a joint probability distribution P^{\max} over the variables $R_1^\sigma, \dots, R_n^\sigma$. The weight vector for the information set $\{R_1^\sigma, \dots, R_n^\sigma\}$ is $\langle a_0, 0, \dots, 0, 1-a_0 \rangle$ and so, from (13), substituting ' $1-r$ ' for ' x ', our degree of confidence that the information is correct would have been:

$$(16) \quad P^{\max*}(R_1^\sigma, \dots, R_n^\sigma) = \frac{a_0}{a_0 + (1-a_0)(1-r)^n}.$$

Hence, in this counterfactual state of affairs our confidence boost would have been:

$$(17) \quad b^{\max*}(\{R_1, \dots, R_n\}) = \frac{P^{\max*}(R_1^\sigma, \dots, R_n^\sigma)}{P^{\max}(R_1^\sigma, \dots, R_n^\sigma)}.$$

Since the prior probability $P^{\max}(R_1^\sigma, \dots, R_n^\sigma) = P(R_1, \dots, R_n) = a_0$, the measure of the proportional confidence-boosting property of the information set is:

$$(18) \quad c_r(\{R_1, \dots, R_n\}) = \frac{b(\{R_1, \dots, R_n\})}{b^{\max}(\{R_1, \dots, R_n\})} = \frac{P^*(R_1, \dots, R_n)}{P^{\max*}(R_1^\sigma, \dots, R_n^\sigma)} \\ = \frac{a_0 + (1-a_0)(1-r)^n}{\sum_{i=0}^n a_i(1-r)^i}$$

This measure is functionally dependent on the expectance measure a_0 and on the reliability measure r . That it is functionally dependent on the expectance measure is desirable, since how much complete overlap there is between the various items of information is relevant to the determination of coherence. But it is unwelcome that this measure is dependent on the reliability measure. Clearly, our pre-theoretical notion of the coherence of an information set does not encompass the reliability of the sources that provide us with its content. So how can we assess the relative coherence of two information sets by this measure of the proportional confidence-boosting property of the information set?

The measure c_r in (18) permits us to construct a quasi-ordering which is independent of the reliability measure. For some pairs of information sets $\{S, S'\}$, $c_r(S)$ will always be greater than $c_r(S')$, no matter what value we choose for r . In this case, S is more coherent than S' . For other pairs of information sets $\{T, T'\}$, $c_r(T)$ is greater than $c_r(T')$ for some values of r and smaller for other values of r . In this case, there is no fact of the matter whether one or the other information set is

more coherent. We will see that this distinction squares with our willingness to make intuitive judgements about the relative coherence of information sets.

Formally, consider two information sets $S = \{R_1, \dots, R_m\}$ and $S' = \{R'_1, \dots, R'_n\}$ and let P be the joint probability distribution for R_1, \dots, R_m and P' the joint probability distribution for R'_1, \dots, R'_n . We calculate the weight vectors $\langle a_0, \dots, a_m \rangle$ for P and $\langle a'_0, \dots, a'_n \rangle$ for P' and construct the following *difference function*:

$$(19) \quad f_r(S, S') = c_r(S) - c_r(S')$$

$f_r(S, S')$ has the same sign for all values of r ranging over the open interval $(0, 1)$ if and only if the measure $c_r(S)$ is always greater than or is always smaller than the measure $c_r(S')$ for no matter what value of r in this interval. We define a coherence relation:

$$(20) \quad \text{For two information sets } S \text{ and } S', S \succeq S' \text{ if and only if } f_r(S, S') \geq 0 \text{ for all values of } r \in (0, 1).$$

' \succeq ' denotes the binary relation of *being more coherent than or equally coherent as*, defined over information sets. This procedure induces a quasi-ordering over a set of information sets.

If the information sets S and S' are of equal size, then it is also possible to determine whether there exists a coherence ordering over these sets *directly* from the weight vectors $\langle a_0, \dots, a_n \rangle$ and $\langle a'_0, \dots, a'_n \rangle$. One needs to evaluate the conditions under which the sign of the difference function is positive for all values of $r \in (0, 1)$. We have shown in Appendix A that:

$$(21) \quad a'_i/a_i \geq \max(1, a'_0/a_0) \quad \forall i = 1, \dots, n-1$$

is a necessary and sufficient condition for $S \succeq S'$ for $n = 2$ and is a sufficient condition for $S \succeq S'$ for $n > 2$.

This is the more parsimonious statement of the condition. However, it is easier to interpret this condition when stated as a disjunction:

$$(22) \quad (i) \quad a'_0 \leq a_0 \wedge a'_i \geq a_i, \quad \forall i = 1, \dots, n-1, \text{ or,}$$

$$(ii) \quad a'_0 \geq a_0 \wedge a'_i/a_i \geq a'_0/a_0, \quad \forall i = 1, \dots, n-1,$$

is a necessary and sufficient condition for $S \succeq S'$ for $n = 2$ and is a sufficient condition for $S \succeq S'$ for $n > 2$.

It is easy to see that (21) and (22) are equivalent.¹⁰




Figure 2 A diagram for the probability distribution for information pairs

We interpret condition (22). For $n = 2$, consider the probability distribution for $S = \{R_1, R_2\}$ represented by the diagram in Figure 2. There are precisely two ways to decrease the coherence in moving from information sets S to S' .¹¹ First, by shrinking the overlapping area between R_1 and R_2 ($a_0' \leq a_0$) and by expanding the non-overlapping area ($a_1' \geq a_1$); and second, by expanding the overlapping area ($a_0' \geq a_0$), while expanding the non-overlapping area to a greater degree ($a_1'/a_1 \geq a_0'/a_0$). The example of *a corpse in Tokyo* in the next section is meant to show that these conditions are intuitively plausible.




Figure 3 A diagram for the probability distribution for information triplets

¹⁰ Assume (21). Either $\max(1, a_0'/a_0) = 1$ or $\max(1, a_0'/a_0) = a_0'/a_0$. In the former case, it follows from the inequality in (21) that $a_0' \leq a_0$ and $a_i' \geq a_i, \forall i = 1, \dots, n-1$. In the latter case, it follows from the inequality in (21) that $a_0' \geq a_0$ and $a_i'/a_i \geq a_0'/a_0, \forall i = 1, \dots, n-1$. Hence, (22) follows. Assume (22). Suppose (i) holds. From the first conjunct in (i), $\max(1, a_0'/a_0) = 1$ and hence from the second conjunct in (i), $a_i'/a_i \geq \max(1, a_0'/a_0) \forall i = 1, \dots, n-1$. Suppose (ii) holds. From the first conjunct in (ii), $\max(1, a_0'/a_0) = a_0'/a_0$ and hence from the second conjunct in (ii), $a_i'/a_i \geq \max(1, a_0'/a_0) \forall i = 1, \dots, n-1$. Hence, (21) follows.

¹¹ We introduce the convention that ‘decreasing’ stands for *decreasing or not increasing*, ‘shrink-

For $n > 2$, consider the probability space for $S = \{R_1, \dots, R_n\}$ in Figure 3. (22) indicates two ways to decrease the coherence in moving from S to S' : first, by shrinking the area in which there is complete overlap between R_1, \dots, R_n ($a_0' \leq a_0$) and by expanding all the areas in which there is no complete overlap ($a_i' \geq a_i, \forall i = 1, \dots, n - 1$); and second, by expanding the area in which there is complete overlap ($a_0' \geq a_0$) and by expanding all the non-overlapping areas to a greater degree ($a_i'/a_i \geq a_0'/a_0, \forall i = 1, \dots, n - 1$). This is a sufficient but not a necessary condition for $n > 2$. The example of *BonJour's ravens* in the next section shows that it may be possible to order two information sets on grounds of our general method in (20) though not on grounds of the sufficient condition in (22).

If we wish to determine the relative coherence of two information sets S and S' of unequal size, there is no short cut. We need to follow our general method in (20) and examine the sign of $f_r(S, S')$ for all values of $r \in (0, 1)$. The example of *Tweety* in the next section will provide an illustration of the procedure to judge the relative coherence of information sets of unequal size.

6. A corpse in Tokyo, BonJour's ravens and Tweety

Does our analysis yield the correct results for some intuitively obvious cases? We consider a comparison (i) of two information *pairs*, (ii) of two information *triples* and (iii) of two information sets of unequal size and show how our method yields intuitively plausible results. In chapter two of *Bayesian Epistemology* (2003), we also show how Lewis's criterion for identifying coherent information sets (1946, p. 338) and the coherence measures suggested by Shogenji (1999), Olsson (2002, p. 250) and Fitelson (2003) yield counter-intuitive results.

i. Information pairs

Suppose that we are trying to locate a corpse from a murder somewhere in Tokyo. We draw a grid of 100 squares over the map of the city and consider it equally probable that the corpse lies somewhere within each square. We interview two partially and equally reliable witnesses. Suppose witness 1 reports that the corpse is somewhere in squares 50 to 60 and witness 2 reports that the corpse is somewhere in squares 51 to 61.

ing' for *shrinking or not expanding*, and 'expanding' for *expanding or not shrinking*. This convention permits us to state the conditions in (21) in a more parsimonious manner and is analogous to the micro-economic convention to let 'preferring' stand for *weak preference*, that is, for *preferring or being indifferent between* in ordinary language.

Call this situation α and include this information in the information set S^α . For this information set, $a_0^\alpha = .10$ and $a_1^\alpha = .02$.

Let us now consider a different situation in which the reports from the two sources overlap far less. In this alternate situation—call it β —witness 1 reports squares 20 to 55 and witness 2 reports squares 55 to 90. This information is contained in S^β . Compared to the information in S^α , the overlapping area shrinks to $a_0^\beta = .01$ and the non-overlapping area expands to $a_1^\beta = .70$. On condition (22)(i), S^β is less coherent than S^α , since $a_0^\beta = .01 \leq a_0^\alpha = .10$ and $a_1^\beta = .70 \geq a_1^\alpha = .02$.

In a third situation γ , witness 1 reports squares 20 to 61 and witness 2 reports squares 50 to 91. S^γ contains this information. Compared to the information in S^α , the overlapping area expands to $a_0^\gamma = .12$ and the non-overlapping area expands to $a_1^\gamma = .60$. On condition (22)(ii), S^γ is less coherent than S^α , since $a_0^\gamma = .12 \geq a_0^\alpha = .10$ and $a_1^\gamma / a_1^\alpha = 30 \geq 1.2 = a_0^\gamma / a_0^\alpha$.

Now let us consider a pair of situations in which no ordering of the information sets is possible. We are considering information pairs, that is, $n = 2$, and so conditions (21) and (22) provide equivalent necessary and sufficient conditions to order two information pairs, *if there exists an ordering*. In situation δ , witness 1 reports squares 41 to 60 and witness 2 reports squares 51 to 70. So $a_0^\delta = .10$ and $a_1^\delta = .20$. In situation ϵ , witness 1 reports squares 39 to 61 and witness 2 reports squares 50 to 72. So $a_0^\epsilon = .12$ and $a_1^\epsilon = .22$. Is the information set in situation δ more or less coherent than in situation ϵ ? It is more convenient here to invoke condition (21). Notice that $a_1^\epsilon / a_1^\delta = 1.10$ is not greater than or equal to $1.20 = \max(1, a_0^\epsilon / a_0^\delta)$, nor is $a_1^\delta / a_1^\epsilon \approx .91$ greater than or equal to $1 = \max(1, a_0^\delta / a_0^\epsilon)$. Hence neither $S^\delta \geq S^\epsilon$ nor $S^\epsilon \geq S^\delta$ hold true.

This quasi-ordering squares with our intuitive judgements. Without having done any empirical research, we conjecture that most experimental subjects would indeed rank the information set in situation α to be more coherent than the information sets in either situations β or γ . Furthermore, we also conjecture that if one were to impose sufficient pressure on the subjects to judge which of the information sets in situations δ and ϵ is more coherent, we would be left with a split vote.

We have reached these results by applying the special conditions in (21) and (22) for comparing information sets. The same results can be obtained by using the general method in (20). Write down the difference functions as follows for each comparison (that is, let $i = \alpha$ and $j = \beta$, let $i = \alpha$ and $j = \gamma$, and let $i = \delta$ and $j = \epsilon$ in turn):

$$(23) f_r(S^i, S^j) = c_r(S^i) - c_r(S^j) =$$

$$\frac{a_0^i + (1 - a_0^i)(1 - r)^2}{a_0^i + a_1^i(1 - r) + a_2^i(1 - r)^2} - \frac{a_0^j + (1 - a_0^j)(1 - r)^2}{a_0^j + a_1^j(1 - r) + a_2^j(1 - r)^2}$$

As we can see in Figure 4, the functions $f_r(S^\alpha, S^\beta)$ and $f_r(S^\alpha, S^\gamma)$ are positive for all values of $r \in (0, 1)$ —so S^α is more coherent than S^β and S^γ . But $f_r(S^\delta, S^\epsilon)$ is positive for some values and negative for other values of $r \in (0, 1)$ —so there is no coherence ordering over S^δ and S^ϵ .




Figure 4 The difference functions for a corpse in Tokyo

ii. Information triples

We return to BonJour's challenge. There is the more coherent set, $S = \{R_1 = [\text{All ravens are black}], R_2 = [\text{This bird is a raven}], R_3 = [\text{This bird is black}]\}$, and the less coherent set, $S' = \{R_1' = [\text{This chair is brown}], R_2' = [\text{Electrons are negatively charged}], R_3' = [\text{Today is Thursday}]\}$. The challenge is to give an account of the fact that S is more coherent than S' . Let us apply our analysis to this challenge.

What is essential in S is that $R_1 \wedge R_2 \vdash R_3$, so that $P(R_3|R_1, R_2) = 1$. But to construct a joint probability distribution, we need to make some additional assumptions. Let us make assumptions that could plausibly describe the degrees of confidence of an amateur ornithologist who is sampling a population of birds:

- (i) There are four species of birds in the population of interest, ravens being one of them. There is an equal chance of picking a bird from each species: $P(R_2) = 1/4$.
- (ii) The variables R_1 and R_2 , whose values are the propositions R_1 and $\neg R_1$, and R_2 and $\neg R_2$, respectively, are probabilistically independent: learning no more than that a raven was (or was not) picked teaches us nothing at all about whether all ravens are black.
- (iii) We have prior knowledge that birds of the same species *often* have the same colour and black may be an appropriate colour for a raven. Let us set $P(R_1) = 1/4$.
- (iv) There is a one in four chance that a black bird has been picked amongst the non-ravens, whether all ravens are black or not, that is, $P(R_3|\neg R_1, \neg R_2) = P(R_3|R_1, \neg R_2) = 1/4$. Since we know that birds of a single species often share the same colour, there is only a chance of $1/10$ that the bird that was picked happens to be black, given that it is a raven and that it is not the case that all ravens are black, that is, $P(R_3|\neg R_1, R_2) = 1/10$.

These assumptions permit us to construct the joint probability distribution over the variables R_1, R_2, R_3 and to specify the weight vector $\langle a_0, \dots, a_3 \rangle$ (see Figure 5).¹²




Figure 5 A diagram for the probability distribution for the set of dependent propositions in BonJour's ravens

What is essential in information set S' is that the propositional variables are probabilistically independent—for example, learning something about electrons presumably does not teach us anything about what day it is today or about the colour of a chair. Let us suppose that the marginal probabilities of each proposition are $P(R_1') = P(R_2') = P(R_3') = 1/4$. We construct the joint probability distribution for R_1', R_2', R_3' and specify the weight vector $\langle a_0', \dots, a_3' \rangle$ in Figure 6.¹³




Figure 6 A diagram for the probability distribution for the set of independent propositions in BonJour's ravens

¹² Since R_1 and R_2 are probabilistically independent, $P(R_1, R_2, R_3) = P(R_1) P(R_2)P(R_3|R_1, R_2)$ for all values of R_1, R_2, R_3 . The numerical values in Figure 5 can be directly calculated.

¹³ Since R_1', R_2' and R_3' are probabilistically independent, $P(R_1', R_2', R_3') = P(R_1')P(R_2')P(R_3')$ for all values of R_1', R_2', R_3' . The numerical values in Figure 6 can be directly calculated.

The information triples do not pass the *sufficient* condition for the determination of the direction of the coherence ordering in (21).¹⁴ So we need to appeal to our general method and construct the difference function:

$$(24) \quad f_{ravens} = f_r(S, S') = \frac{a_0 + (1-a_0)(1-r)^3}{a_0 + a_1(1-r) + a_2(1-r)^2 + a_3(1-r)^3} - \frac{a_0' + (1-a_0')(1-r)^3}{a_0' + a_1'(1-r) + a_2'(1-r)^2 + a_3'(1-r)^3}$$

We have plotted f_{ravens} in Figure 8. This function is positive for all values of $r \in (0, 1)$. Hence we may conclude that S is more coherent than S' , which is precisely the intuition that BonJour wanted to account for.¹⁵

iii. Information sets of unequal size

Finally, we consider a comparison between an information pair and an information triple. The following example is inspired by the paradigmatic example of non-monotonic reasoning about Tweety the penguin. We are not interested in non-monotonic reasoning here, but merely in the question of the coherence of information sets. Suppose that we come to learn from independent sources that someone's pet Tweety is a bird (B) and that Tweety cannot fly, that is, that Tweety is a ground-dweller (G). Considering what we know about pets, $\{B, G\}$ is highly incoherent information. Aside from the occasional penguin, there are no ground-dwelling birds that qualify as pets, and aside from the occasional bat, there are no flying non-birds that qualify as pets. Later, we receive the new item of information that Tweety is a penguin (P). Our extended information set $S' = \{B, G, P\}$ seems to be much more coherent than $S = \{B, G\}$. So let us see whether our analysis bears out this intuition. We construct a joint probability distribution for B, G and P

¹⁴ Clearly the condition fails for $S' \succeq S$, but it also fails for $S \succeq S'$, since $a_2'/a_2 \approx .94 < 1 = \max(1, .25) = \max(1, a_0'/a_0)$.

¹⁵ It is not always the case that an information triple in which one of the propositions is entailed by the two other propositions is more coherent than an information triple in which the propositions are probabilistically independent. For instance, suppose that R_2 and R_3 are extremely incoherent propositions, that is, the truth of R_2 makes R_3 extremely implausible and *vice versa*, and that R_1 is an extremely implausible proposition which in conjunction with R_2 entails R_3 . Then it can be shown that this set of propositions is not a more coherent set than a set of probabilistically independent propositions. This is not unwelcome, since entailments by themselves should not warrant coherence. Certainly, $\{R_1, R_2, R_3\}$ should not be a coherent set when R_2 and R_3 are inconsistent and R_1 contradicts our background knowledge, although $R_1 \wedge R_2 \vdash R_3$. A judgement to the effect that S is more coherent than S' depends both on logical relationships and background knowledge.

together with the marginalized probability distributions for *B* and *G* in Figure 7.

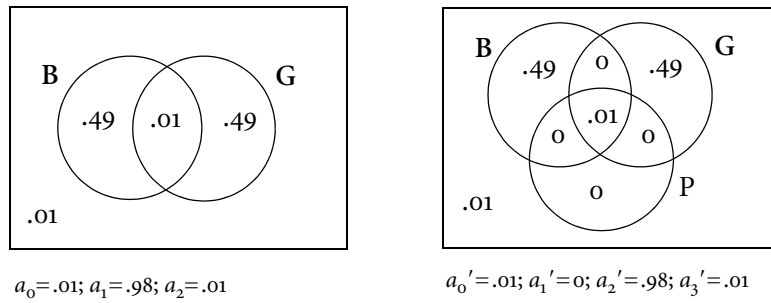


Figure 7 A diagram for the probability distribution for Tweety before and after extension with [Tweety is a penguin]

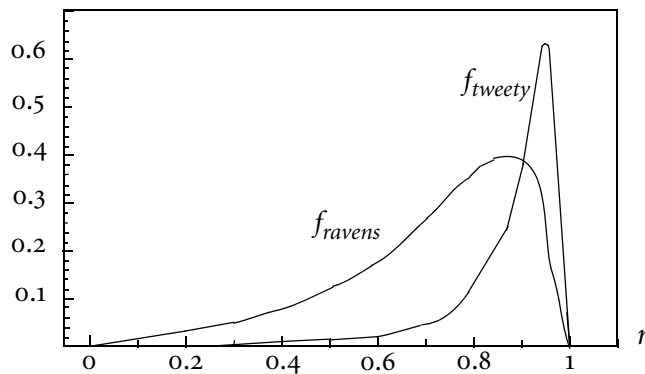


Figure 8 The difference functions for BonJour's ravens and Tweety

Since the information sets are of unequal size, we need to appeal to our general method in (20) and construct the difference function:

$$(25) f_{tweety} = f_r(S', S) =$$

$$\frac{a_0' + (1-a_0')(1-r)^3}{a_0' + a_1'(1-r) + a_2'(1-r)^2 + a_3'(1-r)^3} - \frac{a_0 + (1-a_0)(1-r)^2}{a_0 + a_1(1-r) + a_2(1-r)^2}$$

We have plotted f_{tweety} in Figure 8. This function is positive for all values of $r \in (0, 1)$. We may conclude that S' is more coherent than S , which is precisely the intuition that we wanted to account for.

One might object that on our analysis, the coherence of an information set is dependent on how we partition the information. Consider the information set $S' = \{B, G, P\}$. Suppose that we partition the information as follows: $S'' = \{B \wedge G, P\}$. Since given the background information $B \wedge G$ and P are equivalent propositions, it is easy to show that S'' is a more coherent set than S' . But how is this possible, since the conjunction of the propositions in S' entails the conjunction of the propositions in S'' and vice versa? Note that not only *our* procedure to construct a coherence quasi-ordering but any existing probabilistic coherence measure is subject to this objection.

In response, we claim that the coherence of an information set is subject to how the information is partitioned and information sets with conjunctions of propositions that are equivalent may display different degrees of coherence. If a small percentage of men are unmarried and a small percentage of unmarried people are men in the population, then reports that the culprit is a man and that the culprit is unmarried bring a certain tension to the story. How can that be, we ask ourselves? Aren't most men married and aren't most unmarried people women? The story does not seem to fit together. But if we hear straightaway that the culprit is a bachelor, then this tension is lost. The information that the culprit is a bachelor may be unexpected, since there are so few bachelors. But reporting that the culprit is a bachelor brings no tension to the story. Or consider the following example. There are small settlements of Karaites in Eastern Poland and Lithuania. Though this has been the subject of much controversy, let the Karaites be descendants of East Asian Turkic tribes who accept the (religious) authority of (and only of) the Torah. Suppose that we are told that the culprit was Lithuanian, was a descendant of East Asian Turkic tribes, and accepts the authority of the Torah. Once again, we would be struck by how poorly this information fits together even if we know of Lithuanian Karaites. We are puzzled because of the negative relevance relations between the propositions in question. But if we are told that the culprit is a Karait, is a member of a Lithuanian minority and accepts the authority of the Torah, then we may find this unexpected, but we cannot object that the information does not fit together.

7. Discussion

Equal reliability

We have built into our model the assumption that the sources are equally reliable, that is, that all sources have the same true positive rate p and the same false positive rate q . This seems like an unreasonably strong assumption, since when we are gathering information in the actual world, we typically trust some sources less and some sources more. But our assessment of the relative coherence of information sets has nothing to do with how much we *actually* trust our information sources. As a matter of fact, we may assess the coherence of an information set without having any clue whatsoever who the sources are of the items in this information set or what their degrees of reliability are. An assessment of coherence requires a certain metric that features *hypothetical* sources with certain idealized characteristics. These hypothetical sources are not epistemically perfect, as is usually the case in idealizations. Rather, they are characterized by idealized *imperfections*—their partial reliability. Furthermore, our idealized sources possess the same degree of *internal* reliability and the same degree of *external* reliability. By internal reliability we mean that the sources for each item within an information set are equally reliable, and by external reliability we mean that the sources for each information set are equally reliable.

To see why the same degree of internal reliability is required in our model, consider the following two information sets. Set S contains two equivalent propositions R_1 and R_2 and a third proposition R_3 that is highly negatively relevant with respect to R_1 and R_2 . Set S' contains three propositions R_1' , R_2' and R_3' and every two propositions in S' are just short of being equivalent. One can specify the contents of such information sets such as to make S' intuitively more coherent than S . Our formal analysis will agree with this intuition. Now suppose that it turns out that the actual—that is, the non-idealized—information sources for R_1 , R_1' , R_2 and R_2' are quite reliable and for R_3 and R_3' are close to fully unreliable. We assign certain values to the reliability parameters to reflect this situation and calculate the proportional confidence boosts that actually result for both information sets. Plausible values can be picked for the relevant parameters so that the proportional confidence boost for S actually *exceeds* the proportional confidence boost for S' . This comes about because the actual information sources bring virtually nothing to the propositions R_3 and R_3' and because R_1 and R_2 are indeed equivalent (and hence maximally coherent), whereas R_1' and R_2' are

short of being equivalent (and hence less than maximally coherent). But what we want is an assessment of the relative coherence of $\{R_1, R_2, R_3\}$ and $\{R_1', R_2', R_3'\}$ and not of the relative coherence of $\{R_1, R_2\}$ and $\{R_1', R_2'\}$. The appeal to ideal agents with the same degree of internal reliability in our metric is warranted by the fact that we want to compare the degree of coherence of complete information sets and not of some proper subsets of them. We present a numerical example in Appendix B.

Second, to see why the same degree of external reliability is required in our model, consider some information set S which is not maximally coherent, but clearly more coherent than an information set S' . Our examples in section 6 will do for this purpose. It is always possible to pick two values r and r' so that $c_{r'}(S') > c_r(S)$. To obtain such a result, we need only pick a value of r' in the neighbourhood of 0 or 1 and pick a less extreme value for r , since it is clear from (18) that for r' approaching 0 or 1, $c_{r'}(S')$ approaches 1. This is why coherence needs to be assessed relative to idealized sources that are taken to have the same degree of external reliability.

Indeterminacy

Our analysis has some curious repercussions for the indeterminacy of comparative judgements of coherence. Consider the much-debated problem among Bayesians of how to set the prior probabilities. We have chosen examples in which shared background knowledge (or ignorance) imposes constraints on what prior joint probability distributions are reasonable.¹⁶ In the case of the corpse in Tokyo, one could well imagine coming to the table with no prior knowledge whatsoever about where an object is located in a grid with equal-sized squares. Then it seems reasonable to assume a uniform distribution over the squares in the grid. In the case of BonJour's ravens we modelled a certain lack of ornithological knowledge and let the joint probability distribution respect the logical entailment relation between the propositions in question. In the case of Tweety, one could make use of frequency infor-

¹⁶Note that this is no more than a framework of presentation. Our approach is actually neutral when it comes to interpretations of probability. Following Gillies (2000) and Suppes (2002, Ch. 5) we favour a pluralistic view of interpretations of probability. The notion used in a certain context depends on the application in question. But, if one believes, as a more zealous personalist, that only the Kolmogorov axioms and Bayesian updating impose constraints on what constitute reasonable degrees of confidence, then there will be less room for rational argument and for intersubjective agreement about the relative coherence of information sets. Or, if one believes, as an objectivist, that joint probability distributions can only be meaningful when there is the requisite objective ground, then there will be less occasion for comparative coherence judgements. None of this affects our project. The methodology for the assessment of the coherence of information sets remains the same, no matter what interpretation of probability one embraces.

mation about some population of pets that constitutes the appropriate reference class.

But often we find ourselves in situations without such reasonable constraints. What are we to do then? For instance, what is the probability that the butler was the murderer (B), given that the murder was committed with a kitchen knife (K), that the butler was having an affair with the victim's wife (A), and that the murderer was wearing a butler's jacket (J)? Certainly the prior joint probability distributions over the propositional variables B , K , A , and J may reasonably vary widely for different Bayesian agents and there is little that we can point to in order to adjudicate in this matter. But to say that there is room for legitimate disagreement among Bayesian agents is not to say that anything goes. Certainly we will want the joint probability distributions to respect, among others things, the feature that $P(B|K, A, J) > P(B)$. Sometimes there are enough rational constraints on degrees of confidence to warrant agreement in comparative coherence judgements over information sets. And sometimes there are not. It is perfectly possible for two rational agents to have degrees of confidence that are so different that they are unable to reach agreement about comparative coherence judgements. This is one kind of indeterminacy. Rational argument cannot always bring sufficient precision to degrees of confidence to yield agreement on judgements of coherence.

But what our analysis shows is that this is not the only kind of indeterminacy. Two rational agents may have the same subjective joint probability distribution over the relevant propositional variables and still be unable to make a comparative judgement about two information sets. This is so for situations δ and ϵ in the case of the corpse in Tokyo. Although there is no question about what constitutes the proper joint probability distributions that are associated with the information sets in question, no comparative coherence judgment about S^δ and S^ϵ is possible. This is so because the proportional confidence boost for S^δ exceeds the proportional confidence boost for S^ϵ for some intervals of the reliability parameter, and vice versa for other intervals. If coherence is to be measured by the proportional confidence boost and if it is to be independent of the reliability of the witnesses, then there will not exist a coherence ordering for some pairs of information sets.

In short, indeterminacy about coherence may come about because rationality does not sufficiently constrain the relevant degrees of confidence. In this case, it is our epistemic predicament with respect to the content of the information set that is to blame. However, even when the probabilistic features of a pair of information sets are fully transparent,

it may still fail to be the case that one information set is more coherent than (or equally coherent as) the other. *Prima facie* judgements can be made on both sides, but no judgement *tout court* is warranted. In this case, indeterminacy is not due to our epistemic predicament, but rather, to the probabilistic features of the information sets.

The coherence theory of justification

How does our analysis affect the coherence theory of justification? The coherence theory is meant to be a response to Cartesian scepticism. The Cartesian sceptic claims that we are not justified in believing the story about the world that we have come by from various sources (our senses, witnesses, and so on), since we have no reason to believe that these processes are reliable. There are many variants of the coherence theory of justification. We are interested in versions that hinge on the claim that it is the very coherence of the story of the world that gives us a reason to believe that the story is likely to be true. Can the construction of a coherence quasi-ordering support this claim?

Consider the following analogy. Suppose that we establish that the more a person reads, the more cultured she is, *ceteris paribus*. We conclude from this that if we meet with a very well-read person, then we have a reason to believe that she is cultured. It may not be sufficient reason, but it is a reason nonetheless. Now suppose that we also establish that sometimes no comparison can be made between the amount of reading two people do, since reading comes in many shapes and colours. We can only establish a quasi-ordering over a set of persons according to how well-read they are. This does not stand in the way of our conclusion.

It follows directly from (13) and (18) that

$$(26) \quad P^*(R_1, R_2, \dots, R_n) = \frac{a_0}{a_0 + (1-a_0)(1-r)^n} \times c_r(\{R_1, \dots, R_n\}).$$

Suppose that the measure c_r is greater for an information set S than S' for any value of r . Then S is more coherent than S' . It follows from (26) that the more coherent a particular information set is, the more likely its content is to be true, *ceteris paribus*, in which the *ceteris paribus* clause covers the expectance of the information a_0 and the reliability of the sources r .¹⁷ As in our reasoning about well-read persons, we conclude from this that, if the story of the world is a very coherent infor-

¹⁷The posterior probability that the content of an information set is true is also a function of its size n . We take the size of the information set to be part of its identity conditions and hence this factor does not need to be included in the *ceteris paribus* conditions.

mation set, then we have a reason to believe that its content is likely to be true. Again, it may not be sufficient reason, but it is a reason nonetheless. And similarly, the fact that we can only establish a coherence quasi-ordering over information sets does not stand in the way of this conclusion.

Our claim that the more coherent an information set is, the more likely its content is to be true *ceteris paribus* rests on the assumptions that there is independence between the sources and that we know them to be partially reliable. These assumptions can be relaxed to a certain extent.¹⁸ Our analysis permits us to claim that the coherence of the story about the world provides *some* reason to believe that the story is true, relative to the assumptions, but not that it provides *sufficient* reason. We leave it as an open question whether this is a sufficiently strong claim and whether these are defensible assumptions for the coherence theory of justification to make in providing a successful answer to the Cartesian sceptic.

Coherence and theory choice in science

Where does our analysis leave the claim in philosophy of science that coherence plays a role in theory choice? One can represent a scientific theory T by a set of propositions $\{T_1, \dots, T_m\}$. Let the T_i s be assumptions, scientific laws, specifications of parameters, and so on. It is not plausible to claim that each proposition is independently tested, that is, that each T_i screens off the evidence E_i for this proposition from all other propositions in the theory and all other evidence. The constitutive propositions of a theory are tested in unison. They are arranged into models that combine various propositions in the theory. Different models typically share some of their contents, that is, some propositions in T may play a role in multiple models. It is more plausible to claim that each model M_i is being supported by some set of evidence E_i and that each M_i screens off the evidence E_i in support of the model from the other models in the theory and from other evidence. This is what it means for the models to be supported by independent evidence. There are complex probabilistic relations between the various models in the theory.

Formally, let each M_i for $i = 1, \dots, n$ combine the relevant propositions of a theory T that are necessary to account for the independent

¹⁸ In chapter three of *Bayesian Epistemology*, we show that our analysis remains unaffected when we construct the degree of reliability of the sources as an endogenous variable and Bovens and Olsson (2000) contains a discussion of alternative models of partial reliability in connection with the coherence theory of justification.

evidence E_i for $i = 1, \dots, n$. A theory T can be represented as the union of these M_i s.¹⁹ Let M_i be the variable which ranges over the value M_i stating that all propositions in the model are true and the value $\neg M_i$ stating that at least one proposition in the model is false. In Bayesian confirmation theory, E_i is evidence for M_i if and only if the likelihood ratio

$$(27) \quad x_i = \frac{P(E_i|\neg M_i)}{P(E_i|M_i)} \in (0, 1).$$

Hence, E_i stands to M_i in the same way as $REPR_i$ stands to R_i in our framework. Let us suppose that all the likelihood ratios x_i equal x . We can construct a probability measure P for the constituent models of a theory T and identify the weight vector $\langle a_0, \dots, a_n \rangle$. As in (26),

$$(28) \quad P^*(M_1, \dots, M_n) = \frac{a_0}{a_0 + (1-a_0)x^n} \times c_r(\{M_1, \dots, M_n\}).$$

Suppose that we are faced with two contending theories. The models within each theory are supported by independent items of evidence. Note that the first factor in (28) approximates 1 when the evidence is strong ($x \approx 0$) as well as for large information sets (large n). So, if (i) the evidence for each model is equally strong, as expressed by a single parameter x , and, (ii) either the evidence for each model is relatively strong ($x \approx 0$), or, each theory can be represented by a sufficiently large set of models (large n), then a higher degree of confidence is warranted for the theory that is represented by the more coherent set of models. Of course, we should not forget the caveat that indeterminacy springs from two sources. First, there may be substantial disagreement about the prior joint probability distribution over the variables M_1, \dots, M_n , and second, even in the absence of such disagreement, no comparative coherence judgement may be possible between both theories, represented by their respective constitutive models. But even in the face of our assumptions and the caveats concerning indeterminacy, this is certainly not a trivial result about the role of coherence in theory choice within the framework of Bayesian confirmation theory.²⁰

University of Colorado at Boulder
Department of Philosophy - CB 232
Boulder, CO 80309, USA

LUC BOVENS

¹⁹This account of what a scientific theory is contains elements of both the syntactic view and the semantic view. Scientific theories are characterized by the set of their models, as on the semantic view, and these models (as well as the evidence for the models) are expressed as sets of propositions, as on the syntactic view.

University of Konstanz
Center for Junior Research Fellows - M 682
D-78457 Konstanz, Germany
luc@bovens.org

Department of Philosophy, Logic and
Scientific Method
London School of Economics and Political Science
Houghton Street, London WC2A 2AE, UK

STEPHAN HARTMANN

University of Konstanz
Center for Junior Research Fellows - M 682
D-78457 Konstanz, Germany
s.hartmann@lse.ac.uk

References

- BonJour, L. 1985: *The Structure of Empirical Knowledge*. Cambridge MA: Harvard University Press.
- 1999: 'The Dialectic of Foundationalism and Coherentism' in J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*. Malden MA: Blackwell, pp. 117–42.
- Bovens, L. and E. J. Olsson 2000: 'Coherentism, Reliability and Bayesian Networks'. *Mind*, 109, pp. 685–719.
- Bovens, L. and S. Hartmann 2003: *Bayesian Coherentism*. Oxford: Oxford University Press.
- Dawid, A. P. 1979: 'Conditional Independence in Statistical Theory.' *Journal of the Royal Statistical Society*, ser. B 41, no. 1, pp. 1–31.
- Ewing, A. C. 1934: *Idealism: A Critical Survey*. London: Methuen.
- Fitelson, B. 2003: 'A Probabilistic Theory of Coherence'. *Analysis*, 63, pp. 194–9.
- Gillies, D. 2000: *Philosophical Theories of Probability*. London: Routledge.

²⁰ We are grateful for comments and suggestions from Richard Bradley, Jim Hawthorne, Stephen Leeds, Christian List, Peter Milne, Gerhard Schurz, Josh Snyder, and Timothy Williamson. The University of Konstanz has provided us with a stimulating work environment. Our research was supported by a grant of the National Science Foundation, Science and Technology Studies (SES 00-80580), by the Alexander von Humboldt Foundation (through a Sofja Kovaljevskaja Award, the TransCoop program and the Feodor Lynen Program), the Federal Ministry of Education and Research, and the Program for Investment in the Future (ZIP) of the German Government.

- Kuhn, T. 1977: 'Objectivity, Value Judgment and Theory Choice' in his *The Essential Tension: Selected Studies in Scientific Tradition and Change*. Chicago: University of Chicago Press, pp. 320–9.
- Lewis, C. I. 1946: *An Analysis of Knowledge and Valuation*. LaSalle IL: Open Court.
- Olsson, E. J. 2002: 'What is the Problem of Coherence and Truth?' *Journal of Philosophy*, 94, pp. 246–72.
- Pearl, J. 1988: *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo IL: Morgan Kaufmann.
- Quine, W. V. O. 1960: *Word and Object*. Cambridge MA: MIT Press.
- Reichenbach, H. 1956: *The Direction of Time*. Berkeley: University of California Press.
- Salmon, W. C. 1990: 'Rationality and Objectivity in Science or Tom Kuhn Meets Tom Bayes' in C. Wade Savage (ed.), *Scientific Theories*. Minneapolis: University of Minnesota Press, pp. 175–204.
- 1998: 'Probabilistic Causality' in his *Causality and Explanation*, New York: Oxford University Press, Ch. 14.
- Shogenji, T. 1999: 'Is Coherence Truth-Conducive?' *Analysis*, 59, pp. 338–45.
- Spirtes, P., Glymour, C. and Scheines, R. 2000: *Causation, Prediction, and Search*. 2nd edn. Cambridge MA: MIT Press.
- Suppes, P. 2002: *Representation and Invariance of Scientific Structures*. Stanford: CSLI Publications.

Appendix

A. Proof of Equation (21)

We use the abbreviations $\bar{a}_0 := 1 - a_0$ and $\bar{r} := 1 - r$.

$$\text{Let } c_r(S) = \frac{a_0 + \bar{a}_0 \bar{r}^n}{\sum_{i=0}^n a_i \bar{r}^i}, \quad c_r(S') = \frac{a_0' + \bar{a}_0' \bar{r}^n}{\sum_{i=0}^n a_i' \bar{r}^i}$$

S is more coherent than or equally coherent as S' if and only if

$$\Delta_0 := c_r(S) - c_r(S') \geq 0, \quad \forall r \in (0, 1),$$

Since the denominators of $c_r(S)$ and $c_r(S')$ are greater than 0 for $r \in (0, 1)$, it suffices for our purposes to study when

$$\Delta = \sum_{i=0}^n a_i \bar{r}^i - \sum_{i=0}^n a_i' \bar{r}^i \geq 0.$$

$$\text{Let } D_i := a_0 a_i' - a_0' a_i, \quad \delta_i := a_i' - a_i. \quad (\text{A.1})$$

Note that since $\sum_{i=0}^n a_i = \sum_{i=0}^n a_i' = 1$,

$$\sum_{i=0}^n D_i = a_0 - a_0' = -\delta_0, \quad \sum_{i=0}^n \delta_i = 0, \quad D_0 = 0. \quad (\text{A.2})$$

Using (A.1) and (A.2) one obtains

$$\begin{aligned} \Delta &= \sum_{i=0}^n (a_0 a_i' \bar{r}^i + a_i' \bar{r}^{n+i} - a_0 a_i' \bar{r}^{n+i} - a_0' a_i \bar{r}^i - a_i \bar{r}^{n+i} + a_0' a_i \bar{r}^{n+i}) \\ &= \sum_{i=0}^n [\delta_i \bar{r}^n + D_i (1 - \bar{r}^n)] \bar{r}^i \\ &= \delta_0 \bar{r}^n + \delta_n \bar{r}^{2n} + D_n (\bar{r}^n - \bar{r}^{2n}) + \sum_{i=1}^{n-1} [\delta_i \bar{r}^n + D_i (1 - \bar{r}^n)] \bar{r}^i. \end{aligned}$$

Using the formulae in (A.2) we get $D_n = -\delta_0 - \sum_{i=1}^{n-1} D_i$ and $\delta_n = \delta_0 - \sum_{i=1}^{n-1} \delta_i$; so we obtain after some algebraic manipulations:

$$\Delta = \sum_{i=1}^{n-1} (\bar{r}^i - \bar{r}^n) [\delta_i \bar{r}^n + D_i (1 - \bar{r}^n)]$$

Since $\bar{r}^i - \bar{r}^n > 0$ for $i < n$ and $r \in (0, 1)$, a sufficient (and, for $n=2$, also necessary) condition for S being more coherent than or equally coherent as S' is

$$\delta_i \bar{r}^n + D_i(1 - \bar{r}^n) \geq 0 \quad \forall i=1, \dots, n-1 \text{ and } \forall r \in (0, 1). \quad (\text{A.3})$$

Let $\bar{r}^n =: \lambda \in (0, 1)$ and $f_i(\lambda) = \delta_i \lambda + D_i(1 - \lambda) \quad \forall i=1, \dots, n-1$. Since $f_i(\lambda) \geq 0$ and is monotonic over the range $(0, 1)$,

$$D_i = \lim_{\lambda \searrow 0} f_i(\lambda) \geq 0 \text{ and } \delta_i = \lim_{\lambda \nearrow 0} f_i(\lambda) \geq 0 \quad \forall i=1, \dots, n-1.$$

Hence, (A.3) has the solution

$$D_i \geq 0 \wedge \delta_i \geq 0 \quad \forall i=1, \dots, n-1.$$

We replace D_i and δ_i by the expressions given in (A.1) and obtain:

$$a_i'/a_i \geq a_0'/a_0 \wedge a_i'/a_i > 1 \quad \forall i=1, \dots, n-1.$$

Hence, S is more coherent than or equally coherent as S' if (for $n=2$: if and only if)

$$a_i'/a_i > \max(1, a_0'/a_0) \quad \forall i=1, \dots, n-1.$$

B. Numerical Example

For S , let the joint probability distribution be $P(R_1, R_2, R_3) = .10$, $P(\neg R_1, \neg R_2, R_3) = .80$ and $P(\neg R_1, \neg R_2, \neg R_3) = .10$. For S' , let the joint probability distribution be $P(R_1', R_2', R_3') = .10$, $P(\neg R_1', \neg R_2', R_3') = P(\neg R_1', R_2', \neg R_3') = P(R_1', \neg R_2', \neg R_3') = .01$ and $P(\neg R_1', \neg R_2', \neg R_3') = .87$. Then $\langle a_0, \dots, a_3 \rangle = \langle .10, 0, .80, .10 \rangle$ and $\langle a_0', \dots, a_3' \rangle = \langle .10, 0, .03, .87 \rangle$. From condition (21), S' is more than or equally coherent as S . Now suppose that our information sources for R_1, R_2, R_1' and R_2' are highly reliable, say $p = .75$ and $q = .25$, whereas our information sources for R_3 and R_3' are highly unreliable, say $p^* = .501$ and $q^* = .499$. Then we can use (13) to calculate the posterior joint probability that the information in S and S' , respectively, is true and the posterior joint probability that the information in S and S' , respectively, would have been true had the information been maximally coherent:

$$P^*(R_1, R_2, R_3) = \frac{.10ppp^*}{.10ppp^* + .80qqq^* + .10qqq^*}$$

$$P^{max*}(R_1, R_2, R_3) = \frac{.10ppp^*}{.10ppp^* + .90qqq^*}$$

$$P^*(R_1', R_2', R_3') = \frac{.10ppp^*}{.10ppp^* + .01pqq^* + .01qpq^* + .87qqq^*}$$

$$P^{max^*}(R_1', R_2', R_3') = \frac{.10ppp^*}{.10ppp^* + .90qqq^*}$$

Notice that the normalized confidence boost measure is greater for S than for S' for these assignments of reliability, although, as we have shown above, S' is more coherent than or equally coherent as S.

$$\frac{P^*(R_1, R_2, R_3)}{P^{max^*}(R_1, R_2, R_3)} = .998 > .978 = \frac{P^*(R_1', R_2', R_3')}{P^{max^*}(R_1', R_2', R_3')}$$