Erkenntnis (2005) 63:387–412
DOI 10.1007/s10670-005-4007-z
© Springer 2006

ERIK J. OLSSON

THE IMPOSSIBILITY OF COHERENCE

ABSTRACT. There is an emerging consensus in the literature on probabilistic coherence that such coherence cannot be truth conducive unless the information sources providing the cohering information are individually credible and collectively independent. Furthermore, coherence can at best be truth conducive in a *ceteris paribus* sense. Bovens and Hartmann have argued that there cannot be any measure of coherence that is truth conducive even in this very weak sense. In this paper, I give an alternative impossibility proof. I provide a relatively detailed comparison of the two results, which turn out to be logically unrelated, and argue that my result answers a question raised by Bovens and Hartmann's study. Finally, I discuss the epistemological ramifications of these findings and try to make plausible that a shift to an explanatory framework such as Thagard's is unlikely to turn the impossibility into a possibility.

1. COHERENCE IN EPISTEMOLOGY

In ordinary life we usually rely on the information sources that we have at our disposal, placing our trust in the testimony of other people as well as in that of the senses. Such reliance, as a number of authors have pointed out, is automatic and routine.[1] This is most obvious for the testimony of the senses. Thus, I come to believe that my friend is over there as the direct effect of observing him without in any way inferring his presence from other beliefs I have. But the same is basically true of testimonies from other people. If the secretary tells me that my colleague was in his office just a moment ago, I simply believe it.

While the reception of testimony from various sources is normally unreflective, it is not thereby uncritical. Testimony is accepted so long as there is no explicit reason to doubt the credibility of the reporter, i.e., so long as certain trouble indicators are not present. For example, the information we receive from one source may contradict that received from another or we might have reasons to question the motives of our informant. Is she trying to deceive us? Even an informant with the best of intentions may turn out not to be trustworthy if there are signs that she acquired her information under problematic

circumstances (e.g., under bad lighting conditions). If there are no special reasons for caution, the unreflective mechanism of reliance is invoked and one single testimony suffices to settle the matter, at least for the time being.[2]

Coherence becomes relevant once the reliability of our informants is, for some reason, in doubt, so that we are unable to take that which is being reported at face value. In this case it may pay off to listen to more than one source. If the sources cohere or agree to a large extent in their reporting we may conclude that what they say is true, even though this conclusion could not have been reached as the effect of listening to one of the sources only. If, for instance, the first dubious witness to be queried says that John was at the crime scene, the second that John has a gun and the third that John shortly after the robbery transferred a large sum to his bank account, then the striking coherence of the different testimonies would normally make us pretty confident, notwithstanding their individual dubiousness, that John is to be held responsible for the act.

Lewis made the same point when he asked us to consider a case of "relatively unreliable witnesses who independently tell the same circumstantial story" (1946, p. 246).[3]

> For any one of these reports, taken singly, the extent to which it confirms what is reported may be slight. And antecedently, the probability of what is reported may also be small. But congruence of the reports establishes a high probability of what they agree upon.

The resulting probability of what is agreed need not merely be high but may even suffice for practical certainty:

> Take the case of the unreliable observers who agree in what they report. In spite of the antecedent improbability of any item of such report, when taken separately, it may become practically certain, in a favourable case, merely through congruent relations to other such items, which would be similarly improbable when separately considered.

As Lewis makes clear, the foregoing remarks apply not only to witness reports but quite generally to "evidence having the character of 'reports' of one kind or other – reports of the senses, reports of memory, reports of other persons" (p. 347). Take, for instance, memory reports:

> [S]omething I seem to remember as happening to me at the age of five may be of small credibility; but if a sufficient number of such seeming recollections hang together sufficiently well and are not incongruent with any other evidence, then it may become highly probable that what I recollect is fact. It becomes thus probable just in measure as this congruence would be unlikely on any other supposition which is plausible (p. 352).

Throughout, I will take "testimony" in the widest possible sense to include not only witness testimony but also, for instance, the "testi-

mony of the senses'' and the ''testimony of memory''.[4] Thus, I use ''testimony'' in the same sense in which Lewis uses ''report''.

The foregoing remarks are intended to highlight the *normal* use of coherence, i.e., its employment in inquiries characterized by (1) some of the warning signs being present making it inadvisable to accept testimonies at face value, but (2) there being nonetheless a substantial body of background assumption upon which we can, in fact, rely. Our background information, which is not in doubt in the context of the given inquiry, may tell us, for instance, that the informants are independent of each other and that they, while falling short of full reliability, are nonetheless to be regarded relatively reliable.

What is especially striking about coherence reasoning is that by combining items of information which are in themselves almost worthless one can arrive at a high probability of what is being reported. Indeed, it is salient how little knowledge of the reporters seems necessary for coherence to guarantee high likelihood of truth. We can, it seems, be *almost* entirely ignorant about the quality of our reporters and still arrive at practical certainty as the effect of observing their agreement. At least, this is what Lewis seems to suggest.

There is but a small step from arguing that coherence works under almost total ignorance to holding that it does so even if we remove ''almost''. If coherence is so successful in coping with context where *very little* is taken for granted, could it not also be invoked where *nothing* is? Hence the *anti-skeptical* use of coherence, i.e., the employment of coherence reasoning in skeptical contexts. These contexts are characterized by everything being called into question, except facts of a mere report character. The allowed reports typically state that a person believes or remembers this or that. The claim, then, is that a person can, using coherence reasoning, legitimately recover her trust in her beliefs or memories from this meager base. We can, it is contended, start off with literally nothing – as the skeptic insists – and yet, upon observing the coherence of our *de facto* memories or beliefs, conclude that those memories or beliefs are highly likely to be true.

Thus we are led to the kind of coherence theory advocated by Lewis and BonJour. Both intend to provide a final validation of our empirical knowledge through the anti-skeptical use of coherence reasoning on initially highly dubious data in the form of mere reports on what we believe or (seem to) remember. Their anti-skeptical theories are partly based on certain claims about what is supposed to be true of witness cases, typically accentuating the supposed success of coherence reasoning in such cases. These claims are then said to apply

equally to various skeptical scenarios. A more recent example in the same vain is Coady's (1992) attempt to provide coherence justification of our trust in the testimony of others.

It is important to see that, in all the main epistemological applications of coherence, we are supposed to be presented with some sort of reports. That the reports are present is taken for granted. What is in doubt is whether the contents of the reports – beliefs contents, memory contents etc. – are true. To think otherwise – to think that coherence is to be applied to mere propositions that do not form the contents of some reports – is to commit what I call the *propositional fallacy*. To the best of my knowledge, there are no epistemological applications of coherence as applied to sets of mere propositions, as opposed to sets of reported propositions.

Following Klein and Warfield (1994, 1996), let us say that a measure of coherence is truth conducive if and only if more coherence implies a higher likelihood of truth. As a consequence of the foregoing observation, while it may be true that coherence is not truth conducive when applied to sets of mere propositions – as indeed Klein and Warfield argue – the philosophical relevance of this observation is highly questionable. The epistemologically relevant question is rather whether coherence is truth conducive when predicated of sets of reported propositions. As I have argued, in collaboration with Bovens, the kind of counterexamples to the truth conduciveness of coherence provided by Klein and Warfield have no bearing on the latter question (Bovens and Olsson, 2002).

As has also been argued in the literature – by myself and others – coherence is not truth conducive in the interesting sense unless the circumstances are, in certain respects, fortunate.[5] Thus we need to assume that the reports are collectively independent and individually credible (at least to some degree). Coherence has no effect on the likelihood of truth if the reporters have fudged their story into agreement. Independent reports that are useless when taken singly remain useless when combined, however mutually coherent their contents may be.

Another insight that is gaining acceptance is that coherence can be truth conducive at best in a *ceteris paribus* sense. The most we can hope for is for more coherence to imply higher likelihood of truth, *other things being equal* (Olsson, 2002a).

Let us say that a measure of coherence is *weakly truth conducive* if it is truth conducive *ceteris paribus* given individual credibility and collective independence. It is an open question in the literature whether there are any coherence measures that are weakly truth

conducive. In this paper, I will attempt to show that there are no (non-trivial) measures of that kind.[6]

2. THE CONCEPT OF COHERENCE

The nature of coherence is very much an open question in the epistemological literature. The lack of a definite account of the central concept has been a constant source of embarrassment for coherence theorists. The aim of this paper is not to make a positive contribution in this direction. On the contrary, I intend to show why attempting to define coherence is futile. But in order to do so we must have some basic idea of what a coherence measure is supposed to be.

To avoid the propositional mistake, it is important to make sure that the notion of coherence not be applied to sets of mere propositions, but to sets of *reported* propositions. More precisely, coherence should be predicated of *testimonial systems*. A testimonial system is a set $\mathbf{S} = \{\langle E_1,H_1\rangle, \langle E_2,H_2\rangle,\ldots, \langle E_n, H_n\rangle\}$ where the E_is and H_is are propositions. Intuitively, E_i is a report to the effect that H_i is true. Thus, H_i can be thought of as the content of the report E_i.

We can define the degree of coherence of a testimonial system in two steps:

Step 1: Define degree of coherence for sequences of propositions.
Step 2: Define the degree of coherence of a testimonial system as the degree of coherence of the sequence of its content propositions.

How could we define the degree of coherence of sequences of propositions?[7] Before we take a look at two recent suggestions, it is natural to consider the following simple measure which equates the degree of coherence of a sequence with its joint probability:

$$C_0(H_1,\ldots,H_n) = P(H_1 \wedge \cdots \wedge H_n)$$

Hence,the more likely it is that the proposition are true together, the higher is their degree of coherence. While this may sound plausible at first, it is easy to see that it does not do justice to our intuitive concept of coherence. Suppose a crime has been committed, leaving us wondering who might have done it. Consider the following reports:

Witness no. 1: "Steve did it"
Witness no. 2: "Steve did it"
Witness no. 3: "Steve, Martin or David did it"
Witness no. 4: "Steve, John or James did it"

Which pair of reports would you consider more coherent – that of the first two witness or that of the last two? Presumably, you would favour the reports by the first two witnesses. In other words, you would consider ⟨"Steve did it", "Steve did it"⟩ to be more coherent than ⟨"Steve, Martin or David did it", "Steve, John or James did it"⟩. After all, the first two reports are in perfect agreement, whereas the latter two are not. The C_0 measure, on the other hand, rules that these two pairs are equally coherent, as the joint probability of the one equals the joint probability of the other.

This suggests that we need to measure the degree to which propositions agree. One way to measure agreement, proposed without endorsement in Olsson (2002a) and independently in Glass (2002), would be:

$$C_1(H_1, \ldots, H_n) = \frac{P(H_1 \wedge \cdots \wedge H_n)}{P(H_1 \vee \cdots \vee H_n)}$$

It is easy to see that this measure assigns a maximum coherence value of 1 in all cases of full agreement.

Another way to quantify the degree of agreement is to divide the joint probability not by the probability of the disjunction (as in C_1) but by the product of the propositions' individual probabilities (Shogenji, 1999):[8]

$$C_2(H_1, \ldots, H_n) = \frac{P(H_1 \wedge \cdots \wedge H_n)}{P(H_1) \times \cdots \times P(H_n)}$$

The following example highlights the difference in outcome between applying C_1 and C_2. This time we focus on the following reports:

Witness no. 1: "Steve did it"
Witness no. 2: "Steve did it"
Witness no. 3: "Steve, Martin or David did it"
Witness no. 4: "Steve, Martin or David did it"

Again we ask ourselves which pair of reports exhibits a higher degree of coherence – the first or the last. In this case, one may come up with different answers depending on how one is reasoning. One may, on the one hand, be inclined to say that the degree of coherence is the same on the ground that they are both cases of full agreement. This is also what C_1 dictates. Alternatively, one may be led to think that the first pair is more coherent since what is agreed here is more specific.

This is also what C_2 rules. Thus, C_1 measures how well propositions agree, whereas C_2 measures how striking or salient the agreement is.

Again, the purpose of the foregoing discussion is not to establish anyone of these measures as the correct measure of coherence. Instead, they are just meant to illustrate what a coherence measure could look like. In general, we will mean by a (probabilistic) *coherence measure* any numerical measure that assigns to each sequence $\langle H_1,..., H_n\rangle$ of propositions a number $C(H_1,...,H_n)$ defined solely in terms of the probability of $H_1,...,H_n$ (and their Boolean combinations) and standard arithmetical operations. Clearly, C_0, C_1 and C_2 are all cases in point.

Given a measure of coherence for propositional sequences, we can now define the degree of coherence of a testimonial system. In accordance with what was said above, the degree of coherence of a testimonial system $\mathbf{S} = \{\langle E_1, H_1\rangle, \langle E_2, H_2\rangle,...,\langle E_n, H_n\rangle\}$ equals by definition the degree of coherence of $\langle H_1, H_2,..., H_n\rangle$. This captures the important idea that coherence is supposed to be a property at the level of report contents. Notation: $C_P(\mathbf{S})$ = the degree of coherence assigned to $\mathbf{S}$ by measure C relative to probability distribution P.

To take an example, let

H_1 = "John was at the crime scene"
H_2 = "John has a gun"
H_3 = "John had a motive"

The following is a testimonial system:

$$\mathbf{S} = \{\langle \text{"Smith says that } H_1\text{"}, H_1\rangle, \langle \text{"Mary says that } H_2\text{"}, H_2\rangle, \langle \text{"Karen says that } H_3\text{"}, H_3\rangle\}$$

Moreover, $C(\mathrm{S}) = C(\langle H_1, H_2, H_3\rangle) = C(H_1, H_2, H_3)$.

It is worth emphasizing that coherence, as conceived here, is not conceptually linked to reliability. Coherence is a phenomenon on the level of contents of reports, whereas reliability concerns the relation between a report and its content, i.e., how good a sign the former is of the latter. This is as it should be. We are supposed to be ignorant of the reliability of our data (memories, beliefs, witness reports,...), and so a notion of coherence that depended on reliability would be of little use in an argument against skepticism.

3. WEAK TRUTH CONDUCIVENESS

We are interested in whether there are coherence measures that are truth conducive in the weak sense. Is a more coherent testimonial system therefore likely to be true – at last in fortunate circumstances and in a *ceteris paribus* sense? It is time to spell out what such weak truth conduciveness really amounts to.

First of all: what is the likelihood of truth (probability) of a testimonial system $\mathbf{S} = \{\langle E_1, H_1\rangle,...,\langle E_n, H_n\rangle\}$? It is tempting to take that likelihood to be $P(H_1,...,H_n)$. This, however, would be quite inaccurate. As we have already noted, it can be assumed that the reports have actually been delivered, i.e., that $E_1,..., E_n$ are all true. Hence, $E_1,..., E_n$ are to be counted as evidence. The Principle of Total Evidence dictates that we should, when computing probabilities, take all available evidence into account. Hence, the *probability* of a testimonial system is $P(\mathbf{S}) = P(H_1,..., H_n/E_1,..., E_n)$, i.e., the joint probability of the contents given the reports.[9] We will sometimes refer to $P(\mathbf{S})$ as the posterior probability and to $P(H_1,..., H_n)$ as the prior probability.

We can now define truth conduciveness in the following manner: a coherence measure C is *truth conducive* if and only if: if $C_P(\mathbf{S}) > C_{P'}(\mathbf{S}')$, then $P(\mathbf{S}) > P'(\mathbf{S}')$. In words: a coherence measure is truth conducive whenever more coherence means higher likelihood, regardless of how probabilities are assigned and regardless of what systems are compared. Why do we allow both the probability distribution and the testimonial system to vary between situations that are compared with respect to their relative degree of coherence? Well, why not? I am not aware of any reasons to keep the probability assessments fixed while varying only the testimonial system. By the same token, there seems to be no argument for fixing the testimonial systems while varying the probabilistic assumptions. In the absence of an argument to the contrary, it seems wise to be as liberal as possible in these regards.

How should we understand the "fortunate circumstances", more precisely? By individual credibility is simply meant positive relevance. Thus report E is credible if it raises the probability of its content H, i.e., if $P(H/E) > P(H)$. By collective independence we mean conditional independence in the standard sense. For the purposes of this paper, it will suffice to have that notion defined for the simple case of two reports E_1 and E_2 reporting the same proposition H. The assumption of conditional independence has two parts, corresponding to assuming H true or H false: $P(E_1/H) = P(E_1/H,E_2)$ and $P(E_1/\neg H) = P(E_1/\neg H, E_2)$. It is generally agreed that this notion of independence is adequate in this context and I refer to the literature for further motivation.[10]

Let us turn to the *ceteris paribus* clause. Why should such a clause be imposed in the first place? Suppose we are presented with two testimonial systems, one more coherent than the other. Then, whatever we mean by coherence more precisely, the less coherent (but consistent) system may still be more probable if its reports are individually more credible. In the limit case, the reports of the less coherent system are fully credible, raising the probability of their contents to 1. But it seems unfair to allow such deviations in individual credibility when evaluating the truth conduciveness of a given coherence measure. It seems that factors which have nothing to do with coherence should be kept fixed, especially if they are apt to influence the probability of a testimonial system. Individual credibility is precisely such a factor.

A more controversial issue is what else should be included in the *ceteris paribus* condition. In particular, should we require that the prior probability of the report contents remains fixed? I have argued outgoing from a general analysis of *ceteris paribus* conditions that it should not.[11] The reason is that the prior probability is, in a sense, not sufficiently separable from the degree of coherence. The C_2 measure, for instance, makes coherence heavily dependent on prior probability. Hence prior probability does not belong to the "other things" and so the *ceteris paribus* clause does not require that it remain equal. Individual credibility, by contrast, is separable from the degree of coherence. Changing the individual credibility does not change the degree of coherence.

4. AN IMPOSSIBILITY THEOREM

We are now in a position to address the main issue: are there any measures of coherence that are truth conducive *ceteris paribus* given independence and individual credibility? I will show that even in the simplest of cases there can be no coherence measure that is truth conducive in this weak sense.

We will consider a case of full agreement between independent reports that are individually credible, while respecting the *ceteris paribus* condition. We will show that there are no informative coherence measures that are truth conducive *ceteris paribus* in such a scenario which I will refer to as a *basic Lewis scenario*. The name is appropriate considering Lewis's reference to relatively unreliable witnesses telling the same story. A number of additional constraints

will be imposed on the probabilities involved. The constraints are borrowed from a model proposed by Bovens et al. (2002). That model was in turn devised as an improvement of the model suggested in Olsson (2002b).[12] The most salient feature of this sort of model is that the reliability profile of the witnesses is, in a sense, incompletely known. The witnesses may be completely reliable (R) or they may be completely unreliable (U), and initially we do not know which possibility holds. An interesting consequence of this sort of model is that, from a certain context-dependent level of prior improbability, the posterior probability will be inversely related to the prior: the lower the prior, the higher the posterior. This feature is exploited in the following.

DEFINITION 1. A *basic Lewis scenario* is a pair $\langle \mathbf{S},\mathbf{P}\rangle$ where $\mathbf{S} = \{\langle E_1, H\rangle, \langle E_2, H\rangle\}$ and $\mathbf{P}$ a class of probability distributions defined on the algebra generated by propositions E_1, E_2, R_1, R_2, U_1, U_2 and H such that $P \in \mathbf{P}$ if and only if:

(i) $P(R_i) + P(U_i) = 1$
(ii) $0 < P(H) < 1$
(iii) $P(E_1/H,R_1) = 1 = P(E_2/H,R_2)$
(iv) $P(E_1/\neg H, R_1) = 0 = P(E_2/\neg H, R_2)$
(v) $P(E_1/H,U_1) = P(H) = P(E_2/H,U_i)$
(vi) $P(E_1/\neg H, U_1) = P(H) = P(E_2/\neg H, U_2)$
(vii) $P(R_i/H) = P(R_i) = P(R_i/\neg H)$
(ix) $P(R_1) = P(R_2) > 0$

It can be shown that basic Lewis scenarios satisfy the conditions of individual credibility and independence.

LEMMA 1. (Theorem 3 in Bovens et al., 2002) Let $\langle \mathbf{S},\mathbf{P}\rangle$ be a basic Lewis scenario. Letting $h = P(H), \bar{h} = P(\neg H)$ and $r = P(R_i)$, then

$$P(H/E_1, E_2) = h^* = \frac{(h + r\bar{h})^2}{h + r^2\bar{h}}$$

LEMMA 2. (Bovens et al., 2002, p. 547) Let $\langle \mathbf{S},\mathbf{P}\rangle$ be a basic Lewis scenario. For all r,h^* as a function of h has a unique global minimum for $h \in$]0,1[which is reached at

$$h_{\min} = \frac{r}{1 + r}$$

By calculating the first derivative one can see that h^* increases (decreases) strictly monotonically for $h > (<) h_{\min}$.

Observation 1: $0 < h^* < 1$
Observation 2: $h^* \to 1$ as $h \to 0$
Observation 3: $h_{\min} \to 0$ as $r \to 0$
Observation 4: $h_{\min} \to \frac{1}{2}$ as $r \to 1$

DEFINITION 2. Let C be a coherence measure. C is *informative* in a basic Lewis scenario $\langle \mathbf{S},\mathbf{P} \rangle$ if and only if there are P, $P' \in \mathbf{P}$ such that $\mathrm{C}_P(\mathbf{S}) \neq \mathrm{C}_{P'}(\mathbf{S})$.

DEFINITION 3. A coherence measure C is *truth conducive ceteris paribus* in a basic Lewis scenario $\langle \mathbf{S},\mathbf{P} \rangle$ if and only if: if $C_P(\mathbf{S}) > C_{P'}(\mathbf{S})$, then $P(\mathbf{S}) > P'(\mathbf{S})$ for all $P,P' \in \mathbf{P}$ such that $P(R_i) = P'(R_i)$.

The stipulation that $P(R_i) = P'(R_i)$ is part of the *ceteris paribus* condition. The other part, concerning independence, is guaranteed already by the fact that we are dealing with Lewis scenarios that, so to speak, have independence built into them.

I will make frequent use in the following of the fact that a probability distribution in $\mathbf{P}$ is uniquely characterized by the probability it assigns to H and R_i. Furthermore, for every pair $\langle r,h \rangle$ there is a probability distribution $P_{r,h}$ in $\mathbf{P}$ such that $P(R_i) = r$ and $P(H) = h$.

Observation 5: $P_{r,h\min(r)}(H/E_1,E_2) \to 0$ as $r \to 0$

THEOREM. There are no informative coherence measures that are truth conducive *ceteris paribus* in a basic Lewis scenario.

PROOF. We will seek to establish that if C is truth conducive *ceteris paribus* in a basic Lewis scenario, then C is not informative in such a scenario. We recall that the degree of coherence of an evidential system $\mathbf{S} = \{\langle E_1, H \rangle, \langle E_2,H \rangle\}$ equals the coherence of the pair $\langle H,H \rangle$. Moreover, if C is a coherence measure then $C(\langle H,H \rangle)$ is defined in terms of the probability of H and its Boolean combinations, as explained in Section 2 above. In other words, $C_P(\langle H,H \rangle) = C(h)$ where $h = P(H)$. From what we just said it is clear that in order to show that C is not informative, in the sense of $C_P(\mathbf{S}) = C_{P'}(\mathbf{S})$ for all $P,P' \in \mathbf{P}$, it suffices to prove that $C(h)$ is constant for all $h \in]0,1[$. We will try to accomplish this in two steps, by first showing that $C(h)$ is constant in $I =]0,1/2[$ and then extending this result to the whole interval $]0,1[$. ■

Suppose that C is not constant in I. Hence, there are $h_1,h_2 \in I$ such that $C(h_1) \neq C(h_2)$. We may assume $h_1 < h_2$.

CASE 1. $C(h_1) > C(h_2)$. By Observation 3, $h_{\min}$ goes to 0 as r goes to 0. Since $h_1 > 0$, it follows that there is a probability of reliability r such that $h_{\min} < h_1$. Consider distributions P_{r,h_1} and P_{r,h_2} in **P**. By Lemma 2, $h_{\min}$ is a unique global minimum and h^* is monotonically decreasing for $h > h_{\min}$. Hence, $P_{r,h_1}(h_1/E_1,E_2) < P_{r,h_2}(h_2/E_1,E_2)$. Hence, C is not truth conducive (see Figure 1).

CASE 2. $C(h_1) > C(h_2)$. By Observation 4, $h_{\min}$ goes to 1/2 as r goes to 1. It follows that there is a probability of reliability r such that $h_2 < h_{\min} < 1/2$. Consider distributions P_{r,h_1} and P_{r,h_2} in **P**. By Lemma 2, $h_{\min}$ is a unique global minimum and h^* is monotonically increasing for $h < h_{\min}$. Hence, $P_{r,h_1}(h_1/E_1,E_2) > P_{r,h_2}(h_2/E_1,E_2)$. It follows that C is not truth conducive (see Figure 2).

What has been shown so far is that, if C is truth conducive, C is constant in I.

We will proceed to show that, if C is truth conducive, then C is constant in $I' = [1/2,1[$ as well. Suppose C is truth conducive but not constant in I'. Since C is truth conducive, $C(h) = c$ for all $h \in I$. Since C is assumed not constant in I', there is an $h \in I'$ such that $C(h) \neq c$.

CASE 1. $C(h) > c$. By Observation 2, $P_{r,h}(H/E_1,E_2)$ goes to 1 as h goes to 0. Since $P_{r,h}(H/E_1,E_2) < 1$, there is a $h' \in I$ such that $P_{r,h'}(H/E_1,E_2) > P_{r,h}(H/E_1,E_2)$, whereas $C(h') = c < C(h)$. This contradicts the assumption of C's truth conduciveness (see Figure 3).

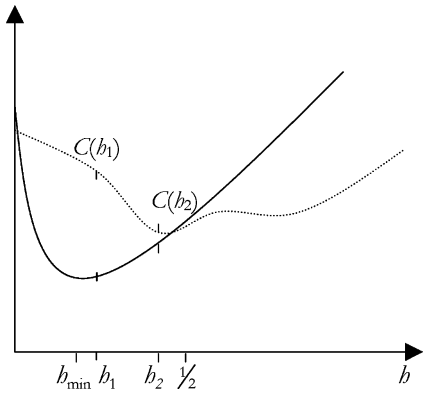

Figure 1. $C(h_1) > C(h_2)$. By choosing r such that $h_{\min} < h_1$ we can construct a counter example to the truth conduciveness of C in the interval $I =]0,1/2[$.

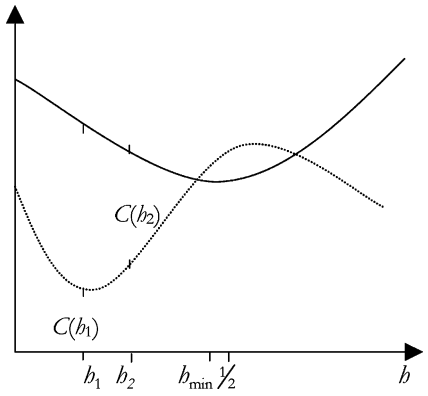

Figure 2. $C(h_1) < C(h_2)$. By choosing r such that $h_{\min} \in]h_2, 1/2[$ we can construct a counter example to the truth conduciveness of C for $h \in]0,1/2[$.

CASE 2. $C(h) < c$. By Observation 5, $P_{r,h\min}(H/E_1,E_2)$ goes to 0 as r goes to 0. By Observation 3, $h_{\min}$ goes to 0 as r goes to 0. It follows by these two observations and the fact that $P(h) > 0$ that there is an r such that $P_{r,h\min}(H/E_1,E_2) < P_{r,h}(H/E_1,E_2)$ with $h_{\min} \in I$. Since $h_{\min} \in I$, $C(h_{\min}) = c > C(h)$. We have shown that there is an h' such that $C(h) < C(h')$ and yet $P_{r,h}(H/E_1, E_2) > P_{r,h'}(H/E_1, E_2)$. Again, we have a clash with the assumption that C is truth conducive (see Figure 4).

We have reached a contradiction and may conclude that, if C is truth conducive, then C is constant not only in I but also in I' so that C is in fact constant in the whole interval $]0,1[$. As we said at the beginning, this is sufficient to establish that, if C is truth conducive *ceteris paribus* for a basic Lewis scenario, then C is not informative in such a scenario QED.

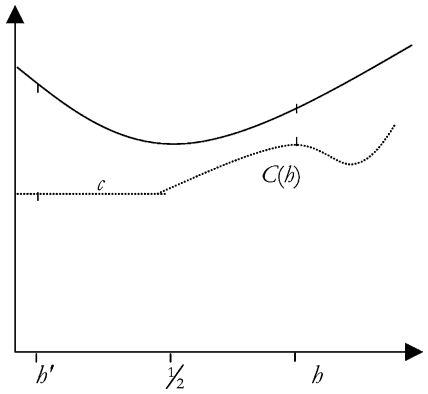

Figure 3. $C(h) > c$. There is then a point h' such that $C(h') = c < C(h)$ but $P_{r,h'}(H/E_1, E_2) > P_{r,h}(H/E_1, E_2)$.

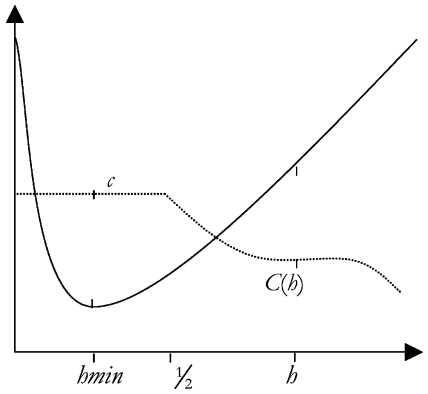

Figure 4. $C(h) < c$. By choosing r so that $P_{r,h\min}(H/E_1,E_2) < P_{r,h}(H/E_1,E_2)$ we get a counter example to the truth conduciveness of C.

5. DISCUSSION

What is it that drives this impossibility theorem? The crucial insight behind it is that exactly how the posterior varies with the prior in this sort of scenario depends not only on the prior probability of what the reports say (H), but also on the prior probability that those reports are reliable (R_i). We can get widely different posteriors depending on how we choose the probability of reliability and, what proves to be absolutely crucial here, the very *kind* of dependence of the prior on the posterior, i.e., the level of improbability at which agreement starts becoming a significant posterior-raising fact is contingent on the prior probability of reliability. In other words, where the curve for the posterior takes on its minimum value is contingent on how we assign the prior probability of reliability. The less probable we take it to be initially that the reports are reliable, the more that minimum will be shifted to the left (referring to the figures above). Thus, not only the absolute but also the *relative* height of the posterior, i.e. what is to count as more or less probable conditional on the evidence, will vary with the initial probability of reliability. Now truth conduciveness, as we have defined it, involves precisely such comparative assessments of posterior probability. Based on this observation alone one might be led to conjecture that there cannot be a non-trivial measure of coherence that is truth conducive in a Lewis type scenario; depending heavily on the precise reliability assumptions, the posterior is underdetermined by facts of coherence alone. What the theorem shows is that this conjecture is correct.

We observe that the degree of coherence in the case we are studying is a function of the prior probability of H. The trick is to show that any such function will, if it is informative, fail to be truth conducive. A counterexample can always be produced by varying the probability of reliability in a strategic way so as to falsify the claim that more coherence, according to the coherence function, implies a higher posterior probability of H. I will comment further on the theorem below in connection with a similar result by Bovens and Hartmann.

The theorem shows that a measure of coherence must pay a price for being truth conducive in a basic Lewis scenario. The price to pay is uninformativeness, i.e., the measure cannot make any distinctions as to coherence but must assign the same coherence value, regardless of the prior probability. The price is substantial since the posterior can vary greatly with the prior in a basic Lewis scenario. Thus, any truth conducive measure of coherence is necessarily useless in the assessment of the height of the posterior. Neither C_0 nor C_2 is uninformative in this sense. Both are heavily dependent on the prior probability of the report contents. This is obvious in the case of C_0 which *is* the prior probability of the report contents, but it is equally true of C_2. As the reader can verify, $C_2(H,H)=\frac{1}{H}$, that is to say, C_2 is inversely related to the prior in a Lewis set-up. Hence, neither C_0 nor C_2 is truth conducive in a basic Lewis scenario. By contrast, the theorem does not rule out the truth conduciveness of C_1. The reason, of course, is that C_1 is uninformative in a basic Lewis scenario, assigning as it does a coherence value of 1 independently of the prior probability of what is being agreed. The generality of the impossibility theorem should be clear. In order to be truth conducive a coherence measure must clearly be truth conducive in a basic Lewis scenario. But this, again, is not possible unless it is uninformative in such a scenario.

6. COMPARISON WITH BOVENS AND HARTMANN'S THEORY

I take the opportunity to comment on a new impossibility result by Bovens and Hartmann (2003) which is similar in spirit to the result proved here. Their result came to my knowledge as I was finalizing this paper, and I regret that I cannot give a full account of their substantial achievements.[13]

The upshot of their reasoning, too, is that it is impossible to define a general truth conducive measure of coherence, that is to say, they

claim to have solved the problem that was first described in Olsson (2002a) as the remaining problem of coherence and truth. Another interesting feature of their book is their proposal for how the coherence theory of justification could be saved from their initial dialectical attack. The main idea is that while it is, in their view, impossible to define a truth conducive measure of coherence in a way that makes all sets comparable with regard to coherence, this is in fact not damaging to the coherence theory. The reason, they say, is that some sets are intuitively not comparable with regard to coherence. If this is true, then the impossibility may be an artifact resulting from attempting to compare what is actually incommensurable. The interesting question is whether there could be truth conducive coherence orderings of sets that are intuitively comparable. In the second chapter of their book, such a "quasi-ordering" is defined with much ingenuity and formal sophistication, and it is argued that it is indeed truth conducive.

One problem with the first two chapters of their book concerns the interpretation of Lewis, whom Bovens and Hartmann rightly refer to as a prominent advocate of the truth conduciveness of coherence. The impossibility result, which is at least initially presented as a problem for Lewis among others, is based on the assumption that the information sources are reliable to a certain, fixed degree which is not subject to change as more information arrives. This sort of reliability is called "exogenous" in the book. Lewis, on the other hand, was quite clear about the fact that in the kind of scenario he took interest in the reliability is initially uncertain and vulnerable to subsequent revision. In fact, I cannot think of any coherence theorist who has shown interest in exogenous reliability. What then, I asked myself, is the philosophical relevance of the impossibility theorem? The problem turned out to be one of presentation only. My question was answered in their third chapter where Bovens and Hartmann proceed to take the more complex situation with uncertain or "endogenous" reliability into account, arguing that their impossibility result can be generalized to cover that sort of case as well.

Here are some remarks on the relation between Bovens and Hartmann's theorem and mine.

Let us start with the question of what they have actually proved. In their argument, they refer to a situation involving three non-equivalent testimonies and two particular assignments of probabilities to the asserted propositions, showing that what probability assignment is associated with a higher posterior probability, i.e. joint probability given the evidence, depends on how reliable the sources

are. Since coherence is assumed to be reliability-independent, there is no coherence measure that makes more coherence imply higher posterior probability in the case at hand. From this, Bovens and Hartmann draw the following conclusion: "we can conclude that there cannot exist a measure of coherence that is probabilistic and induces a coherence ordering for information triples ... and that simultaneously makes it the case that the more coherent the information set, the more confident we are that the information is true, ceteris paribus ..." (p. 21). Yet this does not follow from their example. In fact, any measure that assigns the same coherence value for both probability assignments would do the job. For any such measure, more coherence would imply a higher likelihood of truth for that particular assignment (be it in a trivial sense). Hence, Bovens and Hartmann actually prove a weaker claim than they claim to have proved. What they show is that there is no measure of coherence that makes more coherence imply higher likelihood of truth *ceteris paribus* (under certain fortunate circumstances) *and is informative for the two particular probability assignments figuring in their example*. It is perfectly compatible with their result that there is a totally defined coherence measure that is truth conducive so long as that measure assigns an equal degree of coherence in problematic cases. There is *prima facie* no need for a quasi-ordering. Maybe Bovens and Hartmann should have explored the more conservative strategy of assigning problematic cases an equal degree of coherence before embarking on the more radical path of declaring certain probability assignments to information sets coherentistically incommensurable.

It is worth adding that Bovens and Hartmann have no argument to the effect that it is impossible to define an informative coherence measure that makes more coherence associated with a higher posterior given *any* two distinct assignments of probabilities to an information set. In fact, they show, on p. 23 of their book, that there are assignments for which their way of constructing counterexamples does not work.

The most salient difference between Bovens and Hartmann's result and mine is that I focus on a case of fully agreeing or, more generally, equivalent testimonies. A second difference is that my argument does not hinge on any particular assignment of probabilities to the asserted propositions. Rather, I prove that given *any* two distinct assignments of probabilities to an information set (consisting of two equivalent propositions) it is impossible to define an informative coherence measure that makes more coherence associated with higher posterior (where informativeness is relative to such a basic Lewis

scenario). This suggests, but does not establish, that my way of constructing counterexamples is more powerful than theirs. The trick is to vary the probability of reliability in ways affecting the relationship between posterior and prior probability of what is being asserted, which leads me to the final point of comparison: unlike Bovens and Hartmann's proof, mine is carried out in a scenario that involves endogenous reliability in an essential way; it is crucial to my proof that one can distinguish between two hypotheses regarding the reliability – full reliability and complete unreliability – so that one can meaningfully speak of a "probability of reliability".

It has been suggested to me that Bovens and Hartmann's result is more general than mine. It is true that they have a different conception of what is to be included among the ceteris paribus conditions. In their view, but not in mine, not only the degree of reliability but also the joint prior probability of an information set should be held fixed when comparisons as to coherence are made. Consequently, their aim is to show that coherence is not even truth conducive among information sets having the same prior joint probability, a condition that is satisfied in their example. Had they succeeded in doing this, there might have been a basis for arguing for the greater generality of their result. But, again, what they have proved is actually a weaker statement that involves a reference to the informativeness of the coherence measure relative to a pair of particular probability assignments to an information set (different from the information set figuring in my proof). As far as I can judge, the two results are not only different but indeed logically unrelated.

There is a sense in which my theorem answers an open question raised by Bovens and Hartmann's argument, which does not rule out the existence of measures that are truth conducive in a restricted sense of being correlated with truth for the particularly simple and basic case of equivalent testimonies. My proof shows that there are no such measures that are informative. Indeed, it shows that no such measure exists even if one restricts attention further to sets of equivalent testimonies *of the same size*. The matter, which turns out to be highly non-trivial, depends on the subtle ways in which the choice of a prior probability of reliability influences the relationship between the prior and the posterior probability of what is being asserted.

The urgency of the question of truth conduciveness for equivalent statements depends, at least to some extent, on what kind of application one has in mind. If one is concerned with coherence among witness statements, as legal theorists would typically be, the matter is clearly important because witnesses not only may but often do deliver

equivalent statements. The same is arguably true of most everyday employments of coherence in situations involving possibly unreliable sources. In all these cases, equivalent statements are both possible and frequent. Moreover, the posterior probability can vary tremendously with the prior probability and the number of testimonies, thus making the coherence assessment of the posterior potentially an urgent matter. If we focus our attention on measures that assign the same degree of coherence to equivalent testimonies, as Bovens and Hartmann do, regardless of prior probability and number of testimonies, we have ruled out by definition the possibility of making a coherence assessment of the height of the posterior in such cases. From the point of view of legal and everyday coherence reasoning, there is therefore every reason *not* to focus on measures of this type.

The matter is more delicate if one is exclusively interested in the coherence of one single person's beliefs or memories, as the anti-skeptic typically is. The reason is that the relevant sets whose degrees of coherence are to be ascertained are normally assumed to consist of beliefs having non-equivalent contents. The fact that Bovens and Hartmann's result involves sets of non-equivalent statements makes it somewhat more appealing in this particular context. Having said this, I should remind the reader that, while Lewis and BonJour were both mainly interested in the issue of skepticism, reflection on witness coherence plays an important role in their anti-skeptical argumentation. Both consider witness agreement to be a paradigm case not only of the application of the concept of coherence but also of how coherence, in their view, is positively correlated with likelihood of truth.

Leaving the comparison of the results behind, I remain dissatisfied with the tenor of Bovens and Hartmann's discussion of Cartesian skepticism which conveys the impression that the weak (comparative) truth conduciveness claim upon which they focus is all that is needed for the purposes of a coherence theory of justification; and accordingly that the coherence theorists' sole mistake has been to focus unduly on measures of coherence that impose an ordering, as opposed to a quasi-ordering, on information sets (pp. 26–27). In reality, weak truth conduciveness does not exhaust the coherence theorist's conception of truth conduciveness. Bovens and Hartmann fail to mention that Lewis, for one, was very clear about the need for a more substantial connection between coherence and truth. Lewis thought that we cannot, as a matter of principle, know how reliable our memories are. What we can know is only that they are reliable to some positive degree, though without knowing what that degree is

(Olsson, 2002a; Olsson and Shogenji, 2004). These considerations led him to urge that, for the purposes of a coherence theory, a *high* degree of coherence must be taken to imply a *high* likelihood of truth, regardless of the actual positive degree of reliability of the sources; it is thus insufficient to establish the comparative claim that a *higher* degree of coherence implies a *higher* likelihood of truth. Bovens and Hartmann's introduction of quasi-orderings does little in the direction of establishing the more ambitious contention.

7. COMPARISON WITH THAGARD'S THEORY

An interesting question is whether the impossibility theorem relies in any essential way on the fact that we have been working in a *probabilistic* framework. Is the impossibility result merely an artifact of our modeling assumptions, or does it rather point to a general phenomenon that is independent of particular representations? My view is that the latter is true. The impossibility theorem shows that the likelihood of truth of a given system is seriously underdetermined by facts at the level of propositional contents. In particular, it is underdetermined by facts of coherence.

Thagard's interesting theory of "explanatory coherence" is often seen as a competitor to a probabilistic account of coherence, and it is therefore an interesting question whether it can avoid the problems. This is not the place for an extensive discussion of this rather vexed issue. Rather, I will confine myself to noting, first, that Thagard seems to fall prey to the propositional fallacy. Second, Thagard's meticulous comparison between his own model and the probabilistic setting reveals that these frameworks are in principle very similar – perhaps more so than has been generally appreciated. This gives a *prima facie* reason to believe that a shift to the explanatory framework would not by itself block the impossibility. Third, Thagard himself raises serious doubts as to whether it is possible to measure the degree of coherence of a system, although he does so on grounds that seem different from those upon which our negative conclusion relies.[14]

Let us see how Thagard's theory works. What we begin with is, in the epistemological case, a set of propositions. They can cohere (fit together) or incohere (resist fitting together). Coherence relations include relations of explanation and deduction, whereas incoherence relations include different types of incompatibility, such as logical inconsistency. If two propositions cohere, there is a positive constraint between them. If they incohere, this gives rise to a negative

constraint. The propositions are to be divided into ones that are accepted and ones that are rejected. A positive constraint between two propositions can be satisfied either by accepting both or by rejecting both. Satisfying a negative condition means accepting the one proposition while rejecting the other. A coherence problem, according to Thagard, consists in dividing a set of propositions into accepted and rejected in such a way that the most constraints are satisfied. Thagard presents several different computational models for solving coherence problems, including a model that is based on neural networks.

How different is Thagard's account of coherence from the conception that I have tried to shed light on? As I understand it, coherence is a property of a testimonial system. A testimonial system, we recall, is a set of pairs $\{\langle E_1,H_1\rangle,...,\langle E_n,H_n\rangle\}$ where E_i constitutes testimonial evidence for H_i. The evidence can, for instance, come in the form of testimony from other people, from memory or from the senses. In the Lewis–BonJour tradition, as I reconstruct it, coherence is applied only to structures of this general kind. Lewis, for example, tended to focus on coherence among a person's own memories. It is true that such coherence can raise the probability of other propositions of a purely hypothetical nature, e.g., the hypothesis that the evidence is reliably reported. But this is quite possible without any assessment of the "explanatory coherence" of the hypotheses with the evidence ever taking place. Of course, we could say that in such cases the hypotheses is coherent with the data, and Lewis sometimes adopted this manner of speaking. But I fail to see the point in so doing. Thagard's conception is different from the Lewis–BonJour theory since, in his theory, there are no constraints on what sort of proposition can figure in a coherence problem and hence no restriction on what sets of propositions can "cohere". Sets of propositions in a network will not in general be describable as testimonial systems. Typically, some propositions will have the status of evidence and others the status of (mere) hypotheses that were devised only to explain the evidence. For this reason, it is not clear to me how Thagard avoids the propositional fallacy – if indeed he does avoid it.

Just how different is Thagard's explanationist framework from the probabilistic setting adopted here? The upshot of his admirably detailed comparison of the two frameworks is that it is non-trivial but possible, at least in principle, to translate between the frameworks which is why "it is an open question whether explanationist or probabilist accounts are superior" (p. 271). Given the translatability

between frameworks, one could conjecture that what holds in one framework should hold in the other. In particular, one could conjecture that what is impossible in the probabilistic framework – e.g. defining an interesting measure of coherence that is truth conducive in the weak sense – is just as impossible in Thagard's explanatory framework.

As to my third point, Thagard seems to agree that it is impossible to define a measure of degree of coherence. Having raised the issue, he makes the following observation:

> It would be desirable to define, within the abstract model of coherence as constraint satisfaction, a measure of the degree of coherence of a particular element [with the rest] or of a subset of elements, but it is not clear how to do so. Such coherence is highly nonlinear, since the coherence of an element depends on the coherence of all the elements that constrain it, including elements with which it competes. The coherence of a set of elements is not simply the sum of the weights of the constraints satisfied by accepting them, but depends also on the comparative degree of constraint satisfaction of other elements that negatively constrain them (p. 39).

Thagard goes on to say that his observation cast doubts also on the possibility of quantifying statements such as "Darwin's theory of evolution is more coherent than creationism". Thagard's conclusion is strikingly similar to our own negative results, although the exact relation between his reasons for drawing this conclusion and ours remains an open question.

8. ON THE FUTILITY OF ATTEMPTS TO DEFINE COHERENCE

In a recent paper, BonJour complains that "the precise nature of coherence remains an unsolved problem" (1999, p. 123). He proceeds:

> Spelling out the details of this idea in a way that would allow reasonably precise assessments of comparative coherence, is extremely difficult, at least partly because such an account will dependent on the correct account of a number of more specific and still inadequately understood topics, such as induction, confirmation, probability, explanation and various issues in logic ... (ibid.).

Long before BonJour voiced his doubts, Ewing wrote:

> I think, however, that it is wrong to tie down the advocates of the coherence theory to a precise definition. What they are doing is to describe an ideal that has never yet been completely clarified but is none the less immanent in all our thinking. It would be altogether unreasonable to demand that the moral ideal should be exhaustively defined in a few words, and the same may be true of the ideal of thought. As with the moral ideal, it may well be here that while formulae are helpful, they can provide no

complete stereotyped account, and the only adequate approach is one for which there is no space in this book, namely, a study of what our thought can do at its best by means of numerous examples (1934, p. 231).

Ewing is here suggesting that it might be impossible to capture the concept of coherence in a formula. But what sort of possibility are we talking about here? On a ''weak'' reading, Ewing is saying merely that it would be *practically* impossible, or at least difficult, to state a definition of coherence, as it would require more than ''a few words''. On this interpretation there is little difference in principle between Ewing's view and doubts later raised by BonJour.

Yet, Ewing's remarks also admit a stronger reading. Thus rendered, he is claiming, more radically, that defining coherence is *logically* impossible, that there is no formula or statement, however long, which could do the job adequately. Understood in this way, he is maintaining that there is no systematic account of coherence.

The main result of this paper can be seen as a vindication of Ewing's thesis *under the strong reading*. For it has been shown not only that it is impractical or difficult to define coherence, due to the length of formulas that would be required or to our lack of understanding of crucial notions which such a definition would have to refer to; the impossibility theorem shows that it is outright logically impossible to device such a definition.

There is no mystery about this result. In particular, it does not mean that coherence, while being comprehensible to the human intellect, somehow transcends rational definition. It means simply that the constraints that have been imposed, explicitly or implicitly, on such a definition are jointly incompatible. These constraints include, notoriously, the requirement that a definition of coherence should make that notion, in favourable circumstances, come out as truth conducive *ceteris paribus*. While having coherence imply truth might be too much to ask for, it should at least fall out of a suitable definition that more coherence implies higher probability in a weak *ceteris paribus* sense in favourable circumstances (independence, individual credibility). The constraints also include a condition of informativeness: the degree of coherence should give us some information about how high the posterior is, be it only information about its relative height. The whole point, after all, is to use coherence to assess the likelihood of truth in the face of our supposed initial ignorance about facts of reliability. I have argued that there can be no measure satisfying these requirements. Just as there are no square

circles, there is nothing out there that could play the role coherence is supposed to play. The description of that role is itself incoherent. Small wonder, then, that there has been so little progress in defining coherence.

ACKNOWLEDGEMENTS

I would like to thank Mark Siebel and Staffan Angere for their helpful comments on an earlier version.

NOTES

[1] See, for example, Levi (1991) and Coady (1992). See also Olsson (2003) and Levi (2003).

[2] Cf. Coady (1992), p. 47: "We may have 'no reason to doubt' another's communication even when there is no question of our being gullible; we may simply recognize that the standard warning signs of deceit, confusion, and mistake are not present. This recognition incorporates our knowledge of the witness's competence, of the circumstances surrounding his utterance, of his honesty, of the consistency of the parts of his testimony, and its relation to what others have said, or not said, on the matter."

[3] All references to Lewis concern his 1946 essay *Knowledge and Valuation*.

[4] For an account of different senses of "testimony", including its use in legal contexts, see Chapter 2 in Coady (1992).

[5] Huemer (1997), Bovens and Olsson (2000), Olsson (2001, 2002a) and Olsson and Shogenji (2004).

[6] For a more detailed critique of the Lewis–BonJour tradition of thought, see Olsson (2005).

[7] This section is based on Olsson (2002a).

[8] For a longer discussion of this measure, see Olsson (2001).

[9] Compare Cross (1999).

[10] See Cohen (1977), Jeffrey (1987), Huemer (1997), Bovens and Olsson (2000), Olsson (2002a), Olsson and Shogenji (2004), and Olsson (2005).

[11] See Shogenji (1999) and Olsson (2001). Olsson (2002a) contains a detailed analysis of the *ceteris paribus* clause.

[12] For a discussion of the difference between the two models, see also Olsson (2002c).

[13] Although the book by Bovens and Hartmann which I am here concerned with is listed as published in 2003, it did not in fact appear until mid-2004.

[14] All references to Thagard concern his book from year 2000. For another version of an explanatory coherence theory, see Bartelborth (1996).

REFERENCES

Bartelborth, T.: 1996, *Begründungsstrategien: ein Weg durch die analytische Erkenntnistheorie*, Akademie Verlag, Berlin.

Bonjour, L.: 1985, *The Structure of Empirical Knowledge*, Harvard University Press, Cambridge, Mass.

Bonjour, L.: 1999, 'The Dialectics of Foundationalism and Coherentism', in J. Greco & E. Sosa (eds.), *The Blackwell Guide to Epistemology*, Blackwell, Malden, Mass., 117–142.

Bovens, L., B. Fitelson, S. Hartmann and J. Snyder: 2002, 'Too Odd (not) to Be True: A Reply to Erik J. Olsson', *British Journal for the Philosophy of Science* **53**, 539–563.

Bovens, L. and S. Hartmann: 2003, *Bayesian Epistemology*, Oxford University Press, Oxford.

Bovens, L. and E. J. Olsson: 2000, 'Coherentism, Reliability and Bayesian Networks', *Mind* **109**, 685–719.

Coady, C. A. J.: 1992, *Testimony: A Philosophical Study*, Clarendon Press, Oxford.

Cohen, L. J.: 1977, *The Probable and the Provable*, Clarendon Press, Oxford.

Cross, C. B.: 1999, 'Coherence and Truth Conducive Justification', *Analysis* **59**, 186–193.

Ewing, A. C.: 1934, *Idealism: A Critical Survey*, Methuen, London.

Glass, D. H.: 2002, 'Coherence, Explanation and Bayesian Networks', in *Proceedings of the Irish Conference in AI and Cognitive Science*, Lecture Notes in AI **2646**, Springer, New York, pp. 256–259.

Huemer, M.: 1997, 'Probability and Coherence Justification', *Southern Journal of Philosophy* **35**, 463–472.

Jeffrey, R.: 1987, 'Alias Smith and Jones: The Testimony of the Senses', *Erkenntnis* **26**, 391–399.

Klein, P. and T. A. Warfield: 1994, 'What Price Coherence?', *Analysis* **54**, 129–132.

Klein, P. and T. A. Warfield: 1996, 'No Help For the Coherentist', *Analysis* **56**, 118–121.

Levi, I.: 1991, *The Fixation of Belief and Its Undoing*, Cambridge University Press, Cambridge.

Levi, I.: 2003, 'Contracting from Epistemic Hell Is Routine', *Synthese* **135**, 141–164.

Lewis, C. I.: 1946, *An Analysis of Knowledge and Valuation*, Open Court, LaSalle.

Olsson, E. J.: 2001, 'Why Coherence Is not Truth-Conducive', *Analysis* **61**, 236–241.

Olsson, E. J.: 2002a, 'What Is the Problem of Coherence and Truth?', *The Journal of Philosophy* **99**, 246–272.

Olsson, E. J.: 2002b, 'Corroborating Testimony, Probability and Surprise', *British Journal for the Philosophy of Science* **53**, 273–288.

Olsson, E. J.: 2002c, 'Corroborating Testimony and Ignorance: A Reply to Bovens, Fitelson, Hartmann and Snyder', *British Journal for the Philosophy of Science* **53**, 565–572.

Olsson, E. J.: 2003, 'Avoiding Epistemic Hell: Levion Pragmatism and Inconsistency', *Synthese* **135**, 119–140.

Olsson, E. J.: 2005, *Against Coherence: Truth, Probability, and Justification*, Oxford University Press, Oxford.

Olsson, E. J. and T. Shogenji: 2004, 'Can We Trust Our Memories? C. I. Lewis's Coherence Argument', *Synthese* **142**, 21–41.

Shogenji, T.: 1999, 'Is Coherence Truth-Conducive?', *Analysis* **59**, 338–345.
Thagard, P.: 2000, *Coherence in Thought and Action*, MIT Press, Cambridge, Mass.

Department of Philosophy
Lund University, Kungshuset
SE-222 22, Lund
Sweden
E-mail: Erik_J.Olsson@fil.lu.se