Synthese (2007) 157:347–360
DOI 10.1007/s11229-006-9060-x

ORIGINAL PAPER

On the alleged impossibility of coherence

Wouter Meijs · Igor Douven

Published online: 21 October 2006

Abstract If coherence is to have justificatory status, as some analytical philosophers think it has, it must be truth-conducive, if perhaps only under certain specific conditions. This paper is a critical discussion of some recent arguments that seek to show that under no reasonable conditions can coherence be truth-conducive. More specifically, it considers Bovens and Hartmann's and Olsson's "impossibility results," which attempt to show that coherence cannot possibly be a truth-conducive property. We point to various ways in which the advocates of a coherence theory of justification may attempt to divert the threat of these results.

Keywords Coherence · Truth · Probability · Bovens · Hartmann · Olsson

According to coherentism, a person is justified in holding a belief if, roughly speaking, the belief coheres well with most (or even all, depending on the particular version of coherentism) of her other beliefs, where the notion of coherence is typically circumscribed, at least for starters, in terms of beliefs' hanging well together, or dovetailing with each other, or supporting each other, or in similar metaphorical terms. It is generally thought that a minimal requirement for the tenability of this position is that there be some positive correlation between coherence and truth. Coherence, in other words, should be truth-conducive in the sense that, even if perhaps only under certain specific conditions, one set of beliefs' being more coherent than another entails its being also more probable than that other. In this paper we will call the claim that

W. Meijs
Department of Philosophy, Erasmus University, Rotterdam,
The Netherlands

I. Douven(✉)
Institute of Philosophy, University of Leuven,
Leuven, Belgium
e-mail: igor.douven@hiw.kuleuven.be

coherence is truth-conducive at least in this sense the “Truth-Conduciveness Claim” (TCC).

TCC has faced a number of criticisms, the most recent of which are by Bovens and Hartmann (2003, 2005a, b) and Olsson (2005a, b), who purport to show that it is inconsistent with a number of straightforward and seemingly uncontentious background assumptions; following their usage, we will refer to the arguments they mount to this end as “impossibility results.” These impossibility results are explicitly meant to apply to all definitions of coherence, so that, if correct, it follows that TCC must be abandoned however exactly the notion of coherence is spelled out. Olsson (2005a, p. 4) even goes so far as to conclude that his result “puts us in an excellent position to explain why coherence theorists have been unsuccessful in defining their central notion: coherence is in a sense not definable.” Although these authors are not the first to argue against the truth-conduciveness of coherence, they are to be commended for raising the discussion to a much higher level of precision. Among many other things, they have done pioneering work by pointing the way toward a formal treatment of the key traditional discussions surrounding the notion of coherence. Nevertheless, we do not believe that their impossibility results quite warrant the aforementioned negative claims about coherence. In particular, below we will present five—what we believe to be promising—strategies for a proponent of TCC to divert the threat of the impossibility results.

Previous to that, let us note that Bovens and Hartmann’s and Olsson’s arguments can be conveniently discussed in tandem, for not only do they reach very similar conclusions concerning coherence, their approaches also share many commonalities, two of which it is useful to mention already here. First, both Bovens and Hartmann (henceforth B & H) and Olsson take TCC to pertain not to sets of propositions per se, but to belief systems, where a belief system is a set of propositions that is believed by a particular person (see, e.g., Bovens and Olsson, 2002, p. 140). Second, they propose to analyze belief systems as sets of propositions each of which is supposed to be reported by a witness who is relatively unreliable and independent of all the other witnesses; B & H call such sets “information sets.” Concomitantly, they argue that we should assess TCC with respect to the posterior probability of an information set, that is, the probability that the conjunction of the propositions in the set is true, supposing that we have been informed about each of the propositions by an independent and relatively unreliable witness.

1

Both B & H and Olsson argue that if coherence can be truth-conducive at all, it can only be so in a ceteris paribus (henceforth cp) sense. The most obvious circumstance in which the need for a cp clause may arise is when the witnesses that report the propositions in the information sets are reliable to differing degrees (where the reliability of the witnesses depends in a yet to be specified way on the probability of a report given that a certain proposition is true). For if the witnesses that report the propositions in set S are much more reliable than those that report the propositions in set S', it seems that the posterior probability of S might be larger than that of S' even if S' is intuitively more coherent than S; as long as the reports are not inconsistent we may end up putting a high trust in some rather incoherent information if we believe that the sources are highly unlikely to produce false results. This leads to the conclusion that we should assess TCC only in cases in which the witnesses have equal reliability. According to B & H, we even need a second cp condition, namely equal prior probability (the unconditional probability that the conjunction of the propositions in an information set is true). Olsson does not allow for this cp condition, so that, interestingly, his impossibility result fails to qualify as such on B & H's count. (The acceptance of one or more cp conditions is, in itself, clearly compatible with the truth of TCC, which explicitly leaves open the possibility that only under specific conditions is it true that a more coherent set is more probable.)

According to B & H and Olsson, however, coherence cannot be truth-conducive even in this cp sense. To arrive at this conclusion, they first provide a definition of the reliability of the witnesses. B & H model the witnesses after the example of medical tests: they define the reliability of a witness as a function of her true- and false-positive rates. Olsson, on the other hand, uses a model in which witnesses are either truth-tellers (always tell the truth) or randomizers (randomize between the different options). Like the difference in their views on cp conditions, the difference in their views on how to model witness reliability has some important repercussions for the actual shape of their impossibility results. But both B & H and Olsson show that given their respective explications of witness reliability it is possible that one set of propositions is more probable than another such set if the witnesses are reliable to a degree x_1 and less probable if the witnesses are reliable to a different degree x_2; that is, they show that there exist sets S and S' and values x_1 and x_2 of a reliability parameter such that $p^*_{x_1}(S) > p^*_{x_1}(S')$ and $p^*_{x_2}(S) < p^*_{x_2}(S')$ even if all of the cp conditions that the authors allow are satisfied (here "$p^*_x(S)$" denotes the posterior probability of S given that the witnesses are reliable to degree x). According to B & H and Olsson, altering the reliability of the witnesses should not affect the verdict concerning which of two sets is the more coherent one. And since they show that the posterior probability of one set can be higher than that of another for some values of the reliability parameter, and lower, for other values, they conclude that coherence is not a truth-conducive property.

Their approach is remarkably different from the impossibility result that is proffered in Klein and Warfield (1994). In this paper, Klein and Warfield (henceforth K & W) argue that adding one or more propositions to a set can make it intuitively more coherent, but can never increase its probability. To show that adding propositions to a set can increase its coherence, K & W present a concrete example in which a set S of propositions is claimed to be—and indeed appears to be—intuitively more coherent than one of its subsets S' that contains all but one of S's elements; but of course S cannot be more probable than S' is. Their result differs from the above two impossibility results in two significant ways. First, it only holds good if we conceive of coherence as a property of sets of propositions: as is shown by Bovens and Olsson (2002), adding a belief to a belief system that decreases its prior probability may increase its posterior probability. Secondly, K & W present an intuitive example in which the coherence increases, but the probability does not, while B & H and Olsson give examples in which the posterior probability increases for some values of a reliability parameter and decreases for other values. More specifically, in the results of B & H and Olsson intuitions seem to play only an indirect role (through their assumption that whether a set is more coherent than another does not depend on the reliability of the witnesses).

By itself, it seems advantageous for an impossibility result not to rely too much on intuitions, if only because that makes it harder to escape the result by simply denying the intuitions. However, this kind of consideration only seems relevant when two such results establish the same conclusion. And we believe that this is not the case for the results of, on the one hand, B & H and Olsson and, on the other, K & W. For the former results only seem telling against substantially stronger theses than TCC.

B & H use them to refute the thesis that "if *S* is no less coherent than *S*′, then our degree of confidence that the content of *S* … is true is no less than our degree of confidence that the content of *S*′ is true, ceteris paribus" (2003, p. 11 f). Note, however, that this implies something more than just TCC, namely, that if two sets are equally coherent, then they are also equally (subjectively) probable. And it is unclear why the coherentist ought to commit herself to this extra thesis. To see that their impossibility result leaves TCC unscathed, it is enough to observe that the coherentist is free to declare all combinations of two sets for which it is the case that one of them has a higher posterior probability than the other, given some values of the reliability parameter, and a lower posterior probability than the other, given other values of the reliability parameter, as being equally coherent.[1]

As for Olsson's result, this is only meant to show that there can exist no measures of coherence that satisfy TCC and are informative in basic Lewis scenarios, where a measure of coherence is informative if and only if it can distinguish between two sets with respect to their degree of coherence, and where for present purposes we can characterize a basic Lewis scenario just as one in which two witnesses report the same proposition. Evidently, a measure that makes all basic Lewis scenarios equally coherent for all values of the reliability parameter would make that measure uninformative with respect to these scenarios. To our eyes, however, the proponent of TCC should have no difficulty accepting simply that coherence *is* uninformative in such scenarios. B & H (2003, p. 52) and Fitelson (2003, p. 194) have argued that making sets consisting of equivalent propositions maximally coherent (and, consequently, equally coherent to all other such sets) is a *sine qua non* for any adequate measure of coherence. According to Olsson (2005b), this move would be unsuitable for an application of coherentism to witness statements and, more generally, to all situations involving possibly unreliable sources. In his view, we are often faced with agreeing testimonies, while the posterior probability in such cases can still "vary tremendously with the prior probability and the number of testimonies, thus making the coherence assessment of the posterior probability an urgent matter" (Olsson 2005b). But as TCC is stated, the fact that in some cases the posterior probability may vary while the coherence does not, in no way affects TCC; after all, TCC only applies in cases in which the coherence intuitively increases.[2] Even if one disagrees with B & H's and Fitelson's maximality intuition regarding sets of equivalent propositions—as we do (see Douven and Meijs, 2006a; Meijs, 2006a, b)—one can hardly think that people might have *clear* intuitions about whether one set of equivalent propositions is more (or less) coherent than

[1] This is not the same as arguing that coherence must be a quasi-ordering (which is the conclusion that B & H draw from their impossibility result). For in a quasi coherence ordering, not all sets are ordered with respect to their coherence. But according to the above definition, all pairs of sets can be so ordered. In Meijs and Douven (2005), we have supplied an independent argument to the effect that B & H's account of indeterminate cases can be interpreted as considering all such sets as equally coherent. Therefore, we can see no great obstacles for this move. (In fact, we later found out that in the final version of his 2005b, Olsson makes the same remark with respect to B & H's result.)

[2] This is not necessarily to deny that, as Olsson remarked in personal communication, a conception of coherence that in all cases is sensitive to the posterior probability is preferable to one that is not. Coherence may be put to many uses (cf. Meijs, 2005). However, the important point to note is that such sensitivity is not a feature the proponent of TCC need insist on.

another such set; else it would be hard to explain why so many share B & H's and Fitelson's intuition about sets of equivalent propositions. And if intuitively we do not know how to distinguish such sets with respect to their coherence, it follows that basic Lewis scenarios cannot be used to assess the viability of TCC.

Indeed, it seems that K & W were on the right track with respect to one thing: we can only refute TCC by presenting an example in which one set is intuitively more coherent than another, but at the same time is not (or even cannot be) more probable. But while B & H's impossibility result fails to disown TCC, using their framework it does seem possible to find a counterexample against this thesis. Notice that it is sufficient for this to find any pair of information sets that have equal prior probability for which it is the case that one set is intuitively more coherent than the other, while its posterior probability is not for all values of the reliability parameter higher. The following example of Meijs (2005, p. 99 f), which is a variant of one discussed by B & H (2003, pp. 40–43), fits the said requirements. Consider the set $S = \{C, E, T\}$, with

C: This chair is brown;
E: Electrons are negatively charged;
T: Today is Thursday;

and compare this with $S' = \{B, O, M\}$, with

B: This bird is black;
O: This bird is a crow;
M: This bird has a lifelong mate.

The example assumes that the probability models for these sets are given by Fig. 1. It can easily be checked that the propositions in set S are independent, while the propositions in set S' mutually support each other. Moreover, the propositions in set S' also have a larger relative overlap than those in set S. It thus seems that, by any reasonable criterion, S' should qualify as the more coherent one. Nevertheless, if we consider B & H's measure, it follows that there are values for their reliability parameter such that set S has a higher posterior probability than set S'. Since the prior probabilities in both sets are equal, we believe that this example constitutes a genuine impossibility result for B & H's theory of coherence. Yet we do not believe that this has very important ramifications for TCC. For in order to arrive at such an impossibility result we have to make the same background assumptions that B & H make. And, as we shall argue below, these assumptions are questionable.

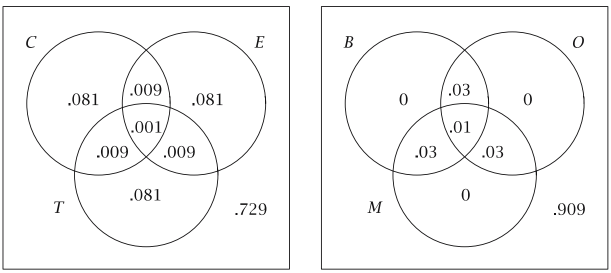

Fig. 1 Diagrams of the probability distributions corresponding to sets S (left) and S' (right)

2

Among the assumptions a proponent of TCC may want to deny are the cp conditions. Recall that according to B & H, we can only assess TCC given equal witness reliability and equal prior probability. These conditions follow from their view on which factors influence our degree of confidence that the reports of a number of different tests are all true. According to them, these factors relate to the trustworthiness of the sources (witness reliability), how expected the results are (prior probability), and how well the results hang together (coherence). This seems a plausible view on the factors that determine our degree of belief.[3] But it is not clear that the only way to take into consideration how expected the results are is through the probability that all of them together are true. We could also measure it by considering how expected the *individual* results are, that is, through the prior probabilities of each of the individual reports. Evidently, for the posterior probability of an information set, it matters not so much what the marginal probabilities of the individual propositions are but much more how probable their conjunction is? But a proponent of TCC could argue that this is precisely the point where considerations of coherence come in: given equal marginal probabilities of the propositions, the coherence of a set goes up as the probability that their conjunction is true increases.

If we accept this suggestion, it follows directly that we should assess TCC under the cp conditions of equal witness reliability and equal *marginal* probabilities of the propositions. In Meijs (2006a), this revised version of the cp conditions is proposed as a solution to a number of counterexamples against B & H's measure of coherence. But it appears that they also hold a promise for resurrecting TCC. For as far as we are aware, there are no examples in the literature on measuring coherence that would count as impossibility results for TCC given these revised cp conditions. (For that we would require two sets with equal marginal probabilities such that one of them is intuitively more coherent than another, but is not more probable for all values of the reliability parameter.) We can provide no general proof that such examples cannot be found, but note that faced with such an example, a proponent of TCC might still be able to hold on to TCC either by rejecting the intuitive clarity of the example or by introducing another cp condition.[4]

Even if no such additional cp condition is necessary, the cp condition of equal marginal probabilities, just as B & H's equal-prior-probabilities condition, would not be admissible according to Olsson's theory of cp conditions. According to Olsson, it only makes sense to hold a property constant when assessing the influence of another property to some phenomenon, if the two properties are capable of independent variation, a condition that is violated if "fixing the value of the one property imposes limitations on the extent to which the other property can consistently vary" (Olsson, 2005a, p. 118). And in his view (*op. cit.*, p. 122 n4), prior probability and coherence cannot be

[3] Note that the view that prior probability is one of the factors that determine the posterior probability of an information set, places us in a good position partially to explain why in Olsson's basic Lewis scenarios the posterior probability can vary. The same fact about prior probability cannot be used to explain the dependence of the posterior probability on the number of propositions—so one would expect that B & H's view that sets consisting of equivalent propositions have maximal coherence, irrespective of the number of propositions, will run into trouble with respect to Lewis scenarios in which sets of different sizes are compared. And indeed, Meijs (2006a) argues that more generally B & H's theory of coherence faces important difficulties when comparing sets of different sizes. In his 2005, Sect. 4.5, Meijs discusses some solutions to these problems.

[4] Especially the condition that both sets should have equal size seems promising in this respect.

independently varied and, consequently, it makes no sense to hold the prior probability fixed when assessing the contribution of coherence to the truth of a belief system. Unfortunately, Olsson does not provide a full-fledged argument for his view on cp conditions and his interpretation of what it means for two properties to be capable of independent variation. Especially his claim that fixing one property (for instance, the prior probability) should place no constraints at all on the range of possible values for the other (for instance, coherence), lacks a decent underpinning and, in effect, strikes us as rather implausible.[5] For instance, consider Shogenji's measure of coherence:

$$C_S\big(\{H_1,\ldots,H_n\}\big) =_{df} \frac{p(H_1 \wedge \cdots \wedge H_n)}{p(H_1) \times \cdots \times p(H_n)}.$$

It can easily be checked that this measure has range $[0,\infty)$, and that setting the prior probability to a value a yields a range $[a,\infty)$, which constitutes a decrease in range for all values of $a > 0$. Nonetheless, we see no reason to conclude from this that C_S and a are incapable of independent variation (an infinite range of values for the degree of coherence remains if we fix the prior probability). Similar remarks apply to holding the marginal probabilities of the propositions constant.[6]

Therefore, we see no compelling reason to accept Olsson's explication of independent variation and, consequently, do not see why one could not attempt to salvage TCC by adopting additional or altogether different cp conditions.[7]

3

As a further strategy, the proponent of TCC may attempt to escape the impossibility results by denying the explication of witness reliability that lies at their basis. As

[5] Olsson bases his view on cp conditions on a discussion by Frazier (1995, p. 119), who stipulates that two properties are incapable of independent variation if it is the case that "when one occurs to a particular degree this is then a certain sign that the other occurs to a particular degree." However, according to Olsson this requirement is too weak, for he feels that in case the other property can take only two other values, this still amounts to "lack of independent variation" (Olsson, 2005a, p. 118). But whereas we would be inclined to agree, we cannot see how this could support Olsson's conclusion that for cp conditions *full* independent variation is required. If there is not much room for independent variation, a ceteris paribus analysis will probably not yield many surprising results, but apart from that we cannot see any injunction against holding a property constant that cannot be fully independently varied from another (for a more elaborate discussion of this point, see Meijs, manuscript in preparation).

[6] But even if we accept Olsson's view on cp conditions, we may still not be committed to the view that strength (or prior probability) and coherence cannot be independently varied. To see this, note that in Douven and Meijs (2006a) a distinction between coherence and incoherence is made. Informally put, a set is any–any coherent if all of its subsets probabilistically support all other subsets in the set, and it is any–any incoherent if all of its subsets probabilistically undermine all other subsets. Now, if we were to make TCC apply only to any–any coherent sets, the problem of lack of independent variation would disappear. As can be easily checked, for all any–any coherent sets, the range of, for instance, Shogenji's measure is $[1,\infty)$, and is thus independent of the prior probability (indeed, for Shogenji's measure the same holds true for all one–any and one–all coherent sets; for definitions, see Douven and Meijs, 2006a). Evidently, such a strategy would constitute a decrease in scope for the coherentist theory, for it would no longer be possible for the coherentist to argue that a set that is less incoherent is more probable than a set that is more incoherent. In how far this decrease is problematic is a matter that deserves more consideration than we can give it here.

[7] Although some of these will be either highly implausible or indeed incapable of sufficient independent variation. For an example of the latter kind one could think of a combination of holding both the prior probabilities and the marginal probabilities constant.

intimated earlier, for the derivation of their impossibility results, B & H and Olsson conceive of witness reliability along different lines. It is noteworthy that there are still other—*prima facie* plausible—ways in which this notion could be defined. For instance, the reliability of a witness could be defined simply as the true-positive rate. And as Shogenji (2006) has noted, given that definition it becomes an open question whether any impossibility results can still be obtained. Consequently, a proponent of TCC may argue that neither B & H nor Olsson has convincingly shown that the impossibility results are not mere artifacts of the definition of witness reliability chosen.[8] To see this more clearly, note that only by first agreeing on some definition of witness reliability can we arrive at an expression for the posterior probability, and without such an expression no impossibility result (at least none following the basic approach sketched above) can be derived. Thus, in order to arrive at a conclusive impossibility result, we would either need a definitive account of how we should measure the reliability of a witness or a general argument that all definitions of witness reliability that satisfy a number of self-evident (or even just plausible) constraints necessarily lead to the existence of impossibility results.

While the first two strategies appear to do nothing more than discredit the impossibility results, this third attempt to salvage coherence as a truth-conducive property of information sets can be used for more constructive ends. For it may be possible to define a reliability parameter such that if one set has a higher posterior probability than another for one value of this parameter, it has a higher posterior for all values. In that case, impossibility results of the type presented by B & H and Olsson are no longer possible. One such attempt is carried out in Meijs (manuscript in preparation), and it is shown there that given the cp condition of equal prior probability, such a definition of reliability can indeed be found. Patently, this does not imply that it is also a satisfactory analysis of witness reliability,[9] nor does it imply that we will never get counterintuitive results: it might still be that for some pairs of sets the intuitively more coherent set is less probable than the other one for all values of the reliability parameter. Yet it does appear to show that this third alternative is a promising one.

4

So far, we have followed B & H and Olsson in defining coherence as a property of information sets. It is generally assumed that this is a necessary condition, since—it is believed—K & W have shown that TCC cannot be true if coherence is a property of sets of propositions. Nevertheless, since we anyway need one or more cp conditions to analyze TCC for information sets, it seems well worth the effort to try and find a number of cp conditions for TCC applied to sets of propositions; K & W certainly have not

---

[8] As Shogenji rightly remarks, B & H's (2003, p. 73 ff) argument to this effect is clearly insufficient.

[9] In response to our suggestion that the coherentist may define witness reliability differently, Olsson (personal communication) has suggested that his own definition derives directly from Lewis's and BonJour's informal characterizations of witness reliability. It seems to us, however, that the coherentist is not committed to those informal characterizations. Be this as it may, Meijs (manuscript in preparation) argues that the alternative definition proposed in that paper is even closer to those characterizations.

shown that TCC cannot be correct under restricted conditions.[10] Here, equal prior probability seems useless: if we are required to keep the prior probability constant, then it will make no sense to inquire whether a higher coherence leads to a higher likelihood of truth. But neither equal strength nor equal marginal probabilities of the propositions can be ruled out directly. This is precisely the line taken by Shogenji (1999) in his attempt to show that coherence can be truth-conducive. He argues that in K & W's example the decrease in prior probability in spite of the increase in coherence is not telling against TCC, given the fact that there is a decrease in the "total individual strength" of the set, where this is defined as the product of the marginal probabilities of the elements of the set.[11]

Shogenji argues by means of an intuitive example that the total individual strength matters for the probability of a set. However, we believe that similar to B & H's description of what factors influence the posterior probability of an information set, we may simply attempt to define a number of different factors that determine the probability of a set of propositions. And if we could argue that coherence and the total individual strength (or the marginal probabilities of the propositions, for that matter) together determine this probability, it would make sense to argue that we should hold the total individual strength (or the marginal probabilities) constant when assessing the truth-conduciveness of coherence.

Given equal total individual strength, it follows directly that Shogenji's measure is truth-conducive (see, for instance, Olsson 2005a, p. 120). Nonetheless, we believe that Shogenji's measure faces an altogether different problem, namely that it does not agree fully with our coherence intuitions. Fitelson (2003, 2004) and Douven and Meijs (2006a) have argued that a measure of coherence should take into account the dependencies between all subsets of a set of propositions (i.e., it should measure the any–any coherence of a set; see note 6). And it can be shown that TCC is false for all extant measures that satisfy this desideratum, or at least that this is so if as our only cp condition we have the assumption of equal total individual strength (the proof is given in the Appendix). It hence appears that until now no satisfactory truth-conducive measures of coherence as a property of sets of propositions have been found. But neither are we aware of any general argument to the effect that such a measure could not be found or that there cannot be any other reasonable cp conditions so that at least some of the measures that have been proposed thus far turn out to be truth-conducive, given those cp conditions.

[10] In Olsson's (2005b) view, it is a fallacy to think that coherence might apply to sets of propositions *per se* (as opposed to sets of reported propositions), given that (he contends) there would be no epistemological use for that kind of coherence. This seems quite wrong, however. Imagine a situation in which we are to choose from among a number of theories that all have been proposed as the explanation of some observed phenomenon. Then there may be no intelligible sense in which the propositions making up these theories could be regarded as being reported to us. Yet it seems to make perfectly good sense to ask which of the theories is the most coherent one (and perhaps it also makes perfectly good sense eventually to accept one of the theories on the basis of our coherence judgement, but that is another matter). Measures of the coherence of sets of propositions (*per se*) would seem well suited to help us answer that question. (The suggestion made in this note is not entirely without problems, but none of these are insurmountable; see Douven and Meijs, 2006a and 2006b.)

[11] See Shogenji (1999, p. 342); also Olsson (2005a, p. 119) for what to our eyes is a slightly clearer presentation of the same point.

5

Finally, we would like to offer an altogether different view on the relation between coherence and probability and, concomitantly, coherence and truth. Rather than letting coherence depend on probabilities—as all proposals considered so far do—the new proposal lets probabilities depend on coherence judgements. It seems undeniable that at least often it is much easier to give a coherence judgement about a given set of propositions—even if only in qualitative terms, such as that a given set is very coherent, or not quite so coherent—than to say how probable the set as a whole is (i.e., how probable the conjunction of its elements is). It also seems reasonable to suppose that such judgements play a role in the determination of our prior probabilities.[12] But even if this hypothesis should be descriptively wrong, it may still be that we, or at least some of us, would do good to let our coherence judgements play a role in determining priors. For it may be that coherence is, in a now to be specified sense, truth-conducive.[13]

We may suppose that it makes sense to speak of the contribution coherence considerations make to our probability judgements, meaning that it is relatively clear to those who do (partly) base their probabilities on such considerations what probability they would have assigned to a given conjunction had they *not* taken into account the coherence of the set of conjuncts, and to those who do not base their probabilities on coherence considerations what probability they would have assigned to a conjunction *had* they taken into account the conjuncts' coherence. Perhaps this assumption is an idealization, but if so, it seems no more problematic than many other idealizations that have been proposed in analytic philosophy.[14] Thus say that for any given person there is a degree of belief function $p(\cdot)$ and a related degree of belief function $p_c(\cdot)$ such that $p(R_1 \wedge \cdots \wedge R_n)$ represents the person's degree of belief in $R_1 \wedge \cdots \wedge R_n$ not based on any considerations regarding the coherence of the set

[12] Indeed, a famous experiment by Kahneman, Slovic, and Tversky (1982) seems to give some support to that hypothesis. In the experiment, respondents were asked to rank, from most to least probable, a number of statements about Linda, of whom they had been given a brief description. It turned out that over 80% of the respondents ranked the sentence "Linda is a bank teller and is active in the feminist movement" above "Linda is a bank teller," that is, in patent conflict with probability theory they assigned a greater probability to a conjunction than to one of the conjuncts. This is generally taken to show that humans do badly at probabilistic reasoning. However, our hypothesis about the relationship between coherence and the assignment of priors offers a more charitable interpretation of the experimental result. The explanation is that when people assess the probability of the conjunction, the question of the coherence of the conjuncts comes up (even if perhaps only unconsciously), and their judgement that these conjuncts indeed cohere (at least to some extent) may lead them to assign a probability to the conjunction that exceeds the ones they assign or would assign to the separate conjuncts, for which the question of coherence of course cannot arise. See B & H (2003, pp. 85–88) for a different attempt to interpret the experimental results charitably. See Olsson (2005, p. 445), though, for a critique of that attempt that does not apply to our proposal.

[13] The claim we are making here parallels the claim, made for instance by Lipton (2004), that explanatory considerations are a guide to inference, in the sense that they have a role, and justly so, in the determination of prior probabilities and sometimes also likelihoods. That such a parallel can be drawn may be no coincidence, given the close relationship between coherence and explanation; see McMullin (1996).

[14] To mention a closely related one: in the context of the discussion concerning the *problem of old evidence* some have supposed that it makes sense to speak of the probability one would have assigned to a given hypothesis had one been unaware of a certain piece of evidence for it; see, e.g., Howson (1984, 1985), and Howson and Urbach (1993, p. 403 ff). For an application of this idea in the scientific realism debate, see Douven (2005).

$\{R_1, \ldots, R_n\}$, and $p_c(R_1 \wedge \cdots \wedge R_n)$ her degree of belief in the same proposition but now based on considerations regarding the coherence of the conjuncts. Which of the two is actual and which counterfactual depends on whether the person does or does not take into account coherence considerations in forming degrees of belief, or at least in forming a degree of belief in $R_1 \wedge \cdots \wedge R_n$. Also note that nothing precludes that $p(R_1 \wedge \cdots \wedge R_n) = p_c(R_1 \wedge \cdots \wedge R_n)$; coherence considerations need not always make a difference. Further, let $P(\cdot)$ be some measure of objective probability. We prefer to think of this kind of probability in terms of limiting frequencies, but its exact nature does not really matter here. Then coherence can be said to be truth-conducive for a given domain of propositions $\mathcal{D} = \{R_i\}$ and for a given person S whose actual and counterfactual degrees of belief functions are $p^S(\cdot)$ and $p^S_c(\cdot)$ (in some order) precisely if, for all conjunctions $R_{i_1} \wedge \cdots \wedge R_{i_n}$ with $n \geqslant 2$ and $R_{i_j} \in \mathcal{D}$ for $j \in \{1, \ldots, n\}$, it holds that

$$P\big(R_{i_1} \wedge \cdots \wedge R_{i_n} \mid p^S_c(R_{i_1} \wedge \cdots \wedge R_{1_n}) > p^S(R_{i_1} \wedge \cdots \wedge R_{1_n})\big) > P\big(R_{i_1} \wedge \cdots \wedge R_{i_n} \mid p^S_c(R_{i_1} \wedge \cdots \wedge R_{1_n}) \leqslant p^S(R_{i_1} \wedge \cdots \wedge R_{1_n})\big).$$

We could say that coherence is truth-conducive (in general) iff this holds for all of us, or at least most of us, and for any domain of propositions. It is not hard to define other senses of truth-conduciveness of coherence by varying upon this definition (most obviously by restricting the universal quantifiers in it in various ways).

We have two remarks on the foregoing. First, we must leave it as an open question whether coherence *is* truth-conducive in the above sense, for anyone or any domain. It seems to us that this is not something that can be known a priori, although it is not immediately evident how one should go about investigating the question empirically. But the important thing to note is that coherence *can* be truth-conducive (in the defined sense) even if we do not know, or will perhaps never know, this. Though we are not willing to bet on this, we just might be the kind of creatures, and inhabit the kind of world, such that sets of propositions that are judged to be coherent by us tend to be true.

Secondly, it is worth emphasizing that the foregoing does not show that attempts to define coherence in terms of probabilities are necessarily futile. One could think so, of course, for if our coherence judgements about sets of propositions underly or at least partly determine the probabilities we assign to those propositions, it might seem that these coherence judgements themselves cannot in turn be based on probabilities; they would seem to lack the required input, so to speak. As others have already noted, however, it is often much easier for us to arrive at conditional probabilities than at unconditional ones.[15] To give an extreme example, many will be at a loss if asked to determine a sharp probability for the proposition that there will be a global war in the next two centuries; by contrast, no one will have difficulty assigning a probability to that proposition conditional on the proposition itself. The crucial point to note now is that some of the probabilistic measures of coherence that have been proposed take as input only conditional probabilities (this is for instance true for the measures $\mathscr{C}_|$ and $\mathscr{F}$ given in the Appendix). Clearly, such measures can well play a role in determining the coherence of a set of propositions whose prior probabilities we are unable to determine directly, provided we are able to determine certain probabilistic correlations between the propositions. Our judgement of the coherence of the set,

[15] See, e.g., Hájek (2003).

based on an application of such a measure, may then subsequently help us to assign prior probabilities to the propositions in the set.[16]

Appendix

Let $[S]$ indicate the set of all ordered pairs of non-empty non-overlapping subsets of $S = \{R_1, \ldots, R_n\}$, and let $[\![S]\!]$ denote the number indicating the cardinality of $[S]$.[17] Finally, let d indicate the following generalized version of the well-known difference measure of confirmation:

$$\mathsf{d}\big(\langle S, S'\rangle\big) \;=_{df}\; p\big(\bigwedge S \,\big|\, \bigwedge S'\big) - p\big(\bigwedge S\big),$$

let r and l be the similarly generalized versions of the ratio and likelihood measures, respectively, and let m be a variable for these generalized measures. In Douven and Meijs (2006a) we defined this family of three measures of coherence:

Definition 1 Given a set $S = \{R_1, \ldots, R_n\}$ and an ordering $\big\langle \hat{S}_1, \ldots, \hat{S}_{[\![S]\!]}\big\rangle$ of the members of $[S]$, the *degree of* m*-coherence* of S is given by the function

$$\mathscr{C}_{\mathsf{m}}(S) \;=_{df}\; \frac{\sum_{i=1}^{[\![S]\!]} \mathsf{m}(\hat{S}_i)}{[\![S]\!]},$$

for $\mathsf{m} \in \{\mathsf{d}, \mathsf{r}, \mathsf{l}\}$.

Meijs (2006b) defines a structurally similar measure of confirmation. Let $|S|$ indicate the set of all subsets of S with cardinality greater than 1, and let $\|S\|$ denote the cardinality of $|S|$.[18] Then Meijs's so-called overlap measure, $\mathscr{O}$, is defined thus:

Definition 2 Given a set $S = \{R_1, \ldots, R_n\}$ and an ordering $\big\langle \tilde{S}_1, \ldots, \tilde{S}_{\|S\|}\big\rangle$ of the members of $|S|$, the *degree of overlap coherence* of S is given by the function

$$\mathscr{O}(S) \;=_{df}\; \frac{\sum_{i=1}^{\|S\|} \mathsf{o}(\tilde{S}_i)}{\|S\|},$$

with

$$\mathsf{o}(\tilde{S}_i) \;=_{df}\; \frac{p\big(\bigwedge \tilde{S}_i\big)}{p\big(\bigvee \tilde{S}_i\big)}.$$

And finally there is Fitelson's (2003, 2004) measure of coherence, $\mathscr{F}$, which is structurally identical to the measures defined in Definition 1 except that it uses the following measure of confirmation where the former use, respectively, d, r, and l:[19]

$$\mathsf{k}\big(\langle S, S'\rangle\big) \;=_{df}\; \frac{p(S' \mid S) - p(S' \mid \neg S)}{p(S' \mid S) + p(S' \mid \neg S)}.$$

[16] We are greatly indebted to Erik Olsson for very helpful comments on a draft version of this paper. Christopher von Bülow's comments are also gratefully acknowledged.

[17] One easily verifies that for any set $S = \{R_1, \ldots, R_n\}$ it holds that $[\![S]\!] = \sum_{i=1}^{n-1} \binom{n}{i}(2^{n-i} - 1)$.

[18] For a set $S = \{R_1, \ldots, R_n\}$ it holds that $\|S\| = 2^n - n - 1$.

[19] In fact, the following is the definition for the cases in which S contains contingent propositions only. For reasons of simplicity, we here ignore the clauses for the other cases.

The following proposition shows that none of the above measures is truth-conducive given Shogenji's cp condition of equal total individual strength (from which it directly follows that none of them is truth-conducive given the cp condition of equal marginal probabilities, either):

Proposition *There exist sets* $S = \{R_1, \ldots, R_n\}$ *and* $S' = \{R'_1, \ldots, R'_m\}$ *such that (i)* $\mathscr{C}_{\mathsf{m}}(S) > \mathscr{C}_{\mathsf{m}}(S')$ *for all* $\mathsf{m} \in \{\mathsf{d}, \mathsf{r}, \mathsf{l}\}$, *(ii)* $\mathscr{O}(S) > \mathscr{O}(S')$, *(iii)* $\mathscr{F}(S) > \mathscr{F}(S')$, *and (iv)* $\prod_{1\leqslant i\leqslant n} p(R_i) = \prod_{1\leqslant i\leqslant m} p(R'_i)$, *but* $p\left(\bigwedge S\right) < p\left(\bigwedge S'\right)$.

Proof: By comparing the following probability models.

Model 1:

R_1	R_2	R_3	probability	R_1	R_2	R_3	probability
T	T	T	.000390625	F	T	T	.099609375
T	T	F	.00001	F	T	F	.00001
T	F	T	.00001	F	F	T	.00001
T	F	F	.099609375	F	F	F	.800350625

Model 2:

R'_1	R'_2	R'_3	probability	R'_1	R'_2	R'_3	probability
T	T	T	.001	F	T	T	.009
T	T	F	.00901	F	T	F	.08101
T	F	T	.00901	F	F	T	.08101
T	F	F	.081	F	F	F	.72896

For we have both $\prod_{1\leqslant i\leqslant 3} p(R_i) = \prod_{1\leqslant i\leqslant 3} p(R'_i) = .10002^3$ and

$$\begin{aligned}
\mathscr{C}_{\mathsf{d}}(\{R_1, \ldots, R_3\}) &= .248362 &>& \quad -.00001 \approx \mathscr{C}_{\mathsf{d}}(\{R'_1, \ldots, R'_3\}),\\
\mathscr{C}_{\mathsf{r}}(\{R_1, \ldots, R_3\}) &= 4.93534 &>& \quad .9997 = \mathscr{C}_{\mathsf{r}}(\{R'_1, \ldots, R'_3\}),\\
\mathscr{C}_{\mathsf{l}}(\{R_1, \ldots, R_3\}) &= 7558.56 &>& \quad .999689 = \mathscr{C}_{\mathsf{l}}(\{R'_1, \ldots, R'_3\}),\\
\mathscr{O}(\{R_1, \ldots, R_3\}) &= .251393 &>& \quad .0404154 = \mathscr{O}(\{R'_1, \ldots, R'_3\}),\\
\mathscr{F}(\{R_1, \ldots, R_3\}) &\approx .002776 &>& \quad -.000156 \approx \mathscr{F}(\{R'_1, \ldots, R'_3\}),
\end{aligned}$$

but also $p(R_1 \wedge R_2 \wedge R_3) < p(R'_1 \wedge R'_2 \wedge R'_3)$.

□

References

Bovens, L., & Hartmann, S. (2003). *Bayesian epistemology*. Oxford: Oxford University Press.

Bovens, L., & Hartmann, S. (2005a). An impossibility result for coherence rankings. *Philosophical Studies* (in press).

Bovens, L., & Hartmann, S. (2005b). Why there cannot be a single probabilistic measure of coherence. *Erkenntnis* (in press).

Bovens, L., & Olsson, E.J. (2002). Believing more, risking less: On coherence, truth and non-trivial extensions. *Erkenntnis, 57*, 137–150.

Douven, I. (2005). Evidence, explanation, and the empirical status of scientific realism. *Erkenntnis, 63*, 253–291.
Douven, I., & Meijs, W. (2006a). Measuring coherence. *Synthese* (in press).
Douven, I., & Meijs, W. (2006b). Bootstrap confirmation made quantitative. *Synthese, 49*, 97–132.
Fitelson, B. (2003). A probabilistic theory of coherence. *Analysis, 63*, 194–199.
Fitelson, B. (2004). Probabilistic coherence from a logical point of view, presentation at the London School of Economics, June 2004 (see http://www.fitelson.org/research.htm).
Frazier, R.L. (1995). Moral relevance and *Ceteris paribus* principles. *Ratio, 7*, 113–127.
Hájek, A. (2003). What conditional probability could not be. *Synthese, 137*, 273–323.
Howson, C. (1984). Bayesianism and support by novel facts. *British Journal for the Philosophy of Science, 35*, 245–251.
Howson, C. (1985). Some recent objections to the Bayesian theory of support. *British Journal for the Philosophy of Science, 36*, 305–309.
Howson, C., & Urbach, P. (1993). *Scientific reasoning: The bayesian approach* (2nd ed.). La Salle, IL: Open Court.
Kahneman, D., Slovic, P., & Tversky, A. (Eds.) (1982). *Judgment under uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
Klein, P., & Warfield, T. (1994). What price coherence? *Analysis, 54*, 129–132.
Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). London: Routledge.
McMullin, E. (1996). Epistemic virtue and theory appraisal. In I. Douven, & L. Horsten (Eds.), *Realism in the sciences* (pp. 13–34). Leuven: Leuven University Press.
Meijs, W. (2005). *Probabilistic measures of coherence*. Ph.D. dissertation. Rotterdam: Erasmus University Rotterdam.
Meijs, W. (2006a). A corrective to Bovens and Hartmann's measure of coherence. *Philosophical Studies* (in press).
Meijs, W. (2006b). Coherence as generalized equivalence. *Erkenntnis* (in press).
Meijs, W. (in preparation). Coherence as a truth-conducive property. Manuscript.
Meijs, W., & Douven, I. (2005). Bovens and Hartmann on coherence. *Mind, 114*, 355–363.
Olsson, E.J. (2005a). *Against coherence*. Oxford: Oxford University Press.
Olsson, E.J. (2005b). The impossibility of coherence. *Erkenntnis* (in press).
Olsson, E.J. (2005c). Review of Bovens and Hartmann (2003). *Studia Logica, 81*, 443–446.
Shogenji, T. (1999). Is coherence truth conducive? *Analysis, 59*, 338–345.
Shogenji, T. (2006). Review of Luc Bovens and Stephan Hartmann's *Bayesian Epistemology*. *Theoria* (in press).