Bayes Factors All the Way: Toward a New View of Coherence and Truth

Abstract: A focus on the conjunction of the contents of witness reports and on the coherence of their contents has had negative effects on the epistemic clarity of the Bayesian coherence literature. Whether or not increased coherence of witness reports is correlated with higher confirmation for some H depends upon the hypothesis in question and upon factors concerning the confirmation and independence of the reports, not directly on the positive relevance of the contents to each other. I suggest that Bayesians should shift focus to "coherence for" an hypothesis--that is, to the definition and analysis of cumulative case arguments in which a body of evidence supports some hypothesis that is not restricted to the conjunction of the contents of reports. Such a shift of focus will be valuable for approaching issues such as the problem of the external world which have interested Bayesian coherentists all along.

1. What H are we talking about?

In the large literature on Bayesianism and coherence, there is a persistent shift between, on the one hand, discussing the truth-conduciveness of the coherence of the contents of testimony for the conjunction of those contents and, on the other hand, discussing the truth-conduciveness of the contents of testimony for some other hypothesis H.

Luc Bovens and Stephan Hartmann begin *Bayesian Epistemology* (2003) with interesting and intuitive examples from biblical stories, including Sarah's miraculous pregnancy and the deception of blind Isaac by his son Jacob. In these examples, the focus is squarely on some hypothesis of interest–"Sarah will bear a child" or "The person standing before me is Esau." Bovens and Hartmann discuss the epistemic effect of coherence of evidential contents upon that hypothesis. For example, Sarah hears from Abraham that God has promised them a son. A stranger comes to the tent and makes the same promise. Then, Sarah (they imagine) begins feeling nauseated in the morning. Bovens and Hartmann point out that the coherence of this set of information appears to confirm the hypothesis that Sarah is indeed going to bear a child. They contrast this with the cognitive dissonance Isaac experiences when he hears a voice that sounds like Jacob's but feels hair-covered hands like Esau's. Isaac's doubts about the identity of his visitor, caused by this evidential incoherence, are eventually overcome when Jacob draws near and allows Isaac to smell his clothes, which smell like Esau (because they are Esau's clothes). Convinced by the coherence between the information given by smell and by touch, Isaac blesses the son before him, who happens to be the deceptive Jacob (Bovens and Hartmann 2003, pp. 1-2). In neither example is the hypothesis in question the conjunction of the contents of all the evidential reports. Rather, the hypothesis is some other conclusion which the subject draws on the basis of the evidence taken as a whole.

Yet later in the book, the focus is somewhat different. For the majority of their discussion of coherence, Bovens and Hartmann are concerned with the degree of confidence one should have in the contents of an entire information set. In other words, the H in question is the conjunction of the facts asserted in all the reports (see, e.g., Bovens and Hartmann 2003, pp. 11-16). The tension between the two different concepts of H is evident when Bovens and Hartmann (2003, p. 18) use the example of three witnesses who, respectively, report the following information:

R_1 The culprit was a woman.
R_2 The culprit had a Danish accent.
R_3 The culprit drove a Ford.

They contrast this with a situation in which the witnesses give the following information:

R'_1 The culprit was wearing Coco Chanel shoes.
R'_2 The culprit had a French accent.
R'_3 The culprit drove a Renault.

Their discussion focuses on the posterior probabilities of the respective conjunctions of these propositions, given various assumptions about the reliability of the witnesses who report them and the prior probabilities. But, while they are interested in the fact that the second set seems more coherent than the first set, they do not make the obvious observation that the second set of testimonial contents "hangs together" better than the first because all three propositions R'_1-R'_3 appear to support an hypothesis which is *not* identical to their conjunction–namely, "The culprit was French." There is no similarly salient H, not identical to the conjunction, which leaps to mind as being *prima facie* supported by R_1-R_3.

Sometimes this shifting focus is avoided because the examples are greatly simplified. If all the witnesses testify to exactly the same thing, and if one focuses on exactly what they *all* testify to, then the distinction between the conjunction of what they testify and some H which their testimonies support disappears in practice. This approach appears in Erik Olsson (2005, pp. 24-25, 2013, p. 30) and in Michael Huemer (2007, p. 338). Bovens and Hartmann also discuss a simplified model involving witness reports that affirm a single, identical proposition (2003, pp. 57-58).

Tomoji Shogenji (2007 and 2013, pp. 2526-7) expressly notes this issue and takes the H which is the conjunction of the contents of testimony to be merely one possible H whose confirmation might be considered in connection with coherence. Shogenji's approach suggests that the case of the conjunction of the contents is merely a special case of the wider and more interesting question in the Bayesian coherence literature: Is the unconditional coherence of the contents of testimony truth conducive for some H which the testimonies appear to support?

I concur with Shogenji that we should not restrict our focus to the impact of coherence upon the conjunction of all testimonial contents in a set. I also consider it of paramount importance to make clear precisely what H we are considering, even if it is the conjunction of all testimonial contents, and to bear in mind in discussions of coherence the confirmation the items provide to *that H*, not simply the confirmation each item of testimony provides to its own *individual* content. By clearly defining H and thinking always in terms of the confirmation provided by the evidence (either individually or jointly) to H, we can avoid mistakes in analysis. By allowing our focus to range over hypotheses other than the conjunction of testimonial contents and by considering the strength of a case for those salient hypotheses, we both maintain the original interest and motivation of the Bayesian coherence project and avoid undue pessimism if some of the goals of that project turn out to be unattainable.

2. Bovens's and Hartmann's confusing analysis of the value of maximal coherence

When constructing their own candidate for a measure of coherence, Bovens and Hartmann say,

> We learned in the last chapter that for any particular value of the reliability parameter r, our confidence boost that [the conjunction of the contents of testimonies is true] is much greater when the information comes to us in the form of maximally coherent information rather than in the form of less than maximally coherent information. Our strategy will be to assess the coherence of an information set by measuring the proportion of the confidence boost that we actually receive, relative to the confidence boost that we would have received *had we received this very same information in the form of maximally coherent information*. (Bovens and Hartmann 2003, pp. 32-33, emphasis in original)

In informal terms, a set of reports is maximally coherent, as Bovens and Hartmann use that phrase, when all witnesses (or tests) say exactly the same thing. When a set of reports is maximally coherent, either all the reports are true or all are false; it is not possible for some to be true while others are false (Bovens and Hartmann 2003, p. 32). Bovens and Hartmann propose to measure less-than-maximal coherence by its "distance" from the confirmation that would have been offered by maximally coherent reports.

A "confidence boost" given by a set of evidence to some hypothesis H is defined by Bovens and Hartmann as the ratio of the posterior probability of H, conditional on the set of testimonies, over its prior (Bovens and Hartmann 2003, p. 33). In other words, "confidence boost" as they are using the term is given by the simple ratio form of the *r* measure of confirmation for some hypothesis H and some evidence E

$$Conf_R(H,E) = \frac{P(H|E)}{P(H).}$$

So their assertion is that a set of reports where all witnesses (or tests) say exactly the same thing provides more confirmation, as measured by $Conf_R$, to the H which is the conjunction of the contents of their testimonies than is provided to the conjunction of the contents by a set of testimonies that do not all say exactly the same thing.

Despite the statement that "we have learned" in the previous chapter that maximal coherence, thus defined, is good for the confirmation boost, there is not any detailed argument to that effect in the previous chapter. It appears that the argument for it is found in this passage concerning a pair of tests for determining the location on the human genome of a gene linked to a particular disease:

> It is plausible that the following three factors affect our degree of confidence that *both* tests are providing us with correct data, i.e. that the faulty gene is indeed located somewhere in the overlapping area....(iii) *How coherent is the information?* This time suppose that the only difference is that in one case both tests identify precisely the same relatively narrow area, whereas in the other case each test identifies a broad area for possible locus with an overlap between these areas that coincides with the relatively narrow area in the first case. Then clearly, our degree of confidence that the *locus* of the faulty gene is in the overlapping area will be lower in the latter than in the former case. (Bovens and Hartmann 2003, pp. 10-11, emphasis in original)

The authors take this example to support the intuition that "the more coherent the information is, the greater our degree of confidence that the information is true, ceteris paribus" (Bovens and

Hartmann 2003, p. 11).

Bovens and Hartmann have, however, misidentified the true moral of the example concerning overlapping tests that identify precisely the same narrow area as compared with tests that have only some overlap between their two broader areas. To see where their analysis goes astray, let us first identify the H in question, as enjoined in the first section of this paper. Since the area of overlap is the same in both scenarios (both when the areas identified by the individual tests are broader and when they are narrower), the H is the same. The H, which is in this case the truth of the conjunction of the "assertions" of the two tests, is something like "The faulty gene lies in *this* (relatively narrow) region," where "this" is the same region in both cases.

Let us now apply the other injunction of the previous section–namely, to consider the confirmation offered by the evidence *to the H in question*. In particular, consider the confirmation given by the *individual* tests to H. The important point to note here is that, all else being equal, an individual test actually offers more support to H when the claim it makes is limited to the area named in H than when the area designated by the evidence merely includes the area named in H as one possibility among others.

To see this point more clearly, consider a scenario in which two independent witnesses are examining a lineup of five possible culprits. In the first version of this scenario, imagine that each witness says specifically that the crime was committed by culprit #5. In the second version, hold constant the confirmation the witnesses give to their individual claims, but suppose instead that one witness says that the crime was committed by either culprit #2, #4, or #5 while the second witness says that the crime was committed by either culprit #1, #3, or #5. The H which is the truth of the conjunction of what both witnesses attest to is, in both cases, "The crime was committed by culprit #5." It is easy to see that, if the witnesses are equally likely to tell the truth in their own testimonies in both cases, the testimonies in the first case provide *individually* stronger support *for H* than in the second case. The reason is rather simple: When Witness 1 testifies to a disjunction that merely includes culprit #5 as one possibility, there are ways for the testimony of Witness 1 to be true while H is false, and the same *mutatis mutandis* for Witness 2. But when each witness gives only one culprit as a possibility, there is no way for what he says to be true while H is false. Hence, *ceteris paribus*, the Bayes factor for the testimony of each witness n

$$\frac{P(W_n|H)}{P(W_n|\sim H)}$$

is higher when the testimony provided by the witness is more specific than when it is more vague, because the probability of the testimony given ~H is higher when the testimony is vaguer. (Arguably, the probability of the vague testimony given H is also sometimes lower than the probability of the specific testimony, but one need not assert this to see that the Bayes factor is stronger for H from the more specific testimony, and that is not the main reason for the disparity in the strength of the factor.)

When the relevant independence factors are held constant, and when the prior probability of H is held constant, a case composed of individually stronger Bayes factors for H gives higher cumulative confirmation to H than a case composed of individually weaker Bayes factors for H. This is uncontroversial. So the confirmation of H by the more specific testimonies, under those circumstances, is stronger than from the vaguer testimonies.

The case of the genetic tests is similar to the case of the specific lineup witnesses vs. the vague witnesses: If the genetic tests are equally likely to tell the truth about their own contents in

both scenarios, and if the H is the same in both scenarios, the case that the faulty gene lies in a relatively narrow area of the genome is stronger when each test makes the specific claim that the fault lies in *that area* than when each test makes a vaguer claim that merely includes that area as a possibility. The increased strength of the case, however, is a function of the fact that each item *by itself* offers more support for H, the conjunction of the contents, in the former case than in the latter. It is confusing to describe this as an instance of the value of the maximal coherence of contents *per se*, even though the two happen to come together in this narrowly defined type of situation.

One might have expected this point to be covered by the *ceteris paribus* conditions in Bovens's and Harmann's analysis, since they require test reliability to be held constant. But their reliability parameter concerns not the test's reliability *vis a vis* H but only the test's reliability *vis a vis* its own content, which might or might not be identical to H (Bovens and Hartmann 2003, pp. 16-17). They thus have *no* parameter that indicates the confirmation offered by the individual tests/testimonies *to H* where the contents of the tests/testimonies are not identical to each other, and hence no way of discussing whether such a parameter is, or should be, held constant across scenarios.

In the precise types of scenario envisaged in these examples, it would not be possible to hold the individual confirmations of H equal without violating one or more of Bovens's and Hartmann's other *ceteris paribus* or independence conditions.[1] But once this issue is brought to the fore, we will have to decide whether it is more important to hold constant the confirmation the individual testimonies offer to their individual contents or the confirmation they offer to the conjunction of their contents, and whether this importance might even vary from one type of scenario to another. This issue is obscured by the way that Bovens and Hartmann frame their conclusion--namely, that maximal coherence of the contents is valuable *per se* to confirmation.[2]

3. Other Bayesian reasons for thinking that coherence of contents is not generally truth conducive

The most well-known reasons for believing that coherence of the contents of testimony is not in general truth conducive are, of course, the impossibility results in the literature. Olsson (2005, pp. 134ff, pp. 211ff) has one impossibility result; Bovens and Hartmann (2003, pp. 19ff) have an impossibility result that is similar in some respects to Olsson's and differs in others. (For example, Bovens and Hartmann hold constant the prior probability of the hypothesis while Olsson does not.)

Bovens and Hartmann, however, do not take from their impossibility result the moral that coherence of contents is not truth conducive for the conjunction (2003, pp. 22ff). Instead, they deem that, by definition, in any situation in which the coherence of two sets of testimonial contents really is comparable and their *ceteris paribus* conditions are held constant, the more

[1] Because of the possibility of negative dependence of reports given the negation of the conjunction of the contents, it may be possible to counterexample the claim that reports with identical contents *always* confirm the conjunction of their contents more than reports with non-identical contents, even given all of Bovens's and Hartmann's conditions. I do not claim to offer a counterexample here. Bovens's and Hartmann's own impossibility results (2003, pp. 19-22, 35) show that it is not possible to generalize about the truth-conduciveness of "closeness" to identity of contents.

[2] This section underscores the fact that, as Shogenji has pointed out (2013, pp. 2539-40), treating the conjunction of the contents as the H in question across the board is in some ways epistemically confusing and hence not ideal for analysis.

coherent set will always have the higher posterior probability. This amounts to a stipulatation that, by definition, a coherence ordering between two sets of testimonies exists only where the coherence of the contents is, *ceteris paribus*, a positive factor. Hence, where coherence of contents does not turn out to be helpful to the posterior probability of the conjunction of the contents despite the fact that the *ceteris paribus* conditions are held constant, they conclude that the two sets are actually non-comparable in terms of coherence; neither is more coherent than the other. This has the advantage of guaranteeing that coherence of the contents is truth conducive, but, as others have noted, it risks trivializing the argument and abandoning the pre-analytic concept of coherence in the name of locking in truth-conduciveness. (See Olsson 2005, pp. 153-4, Shogenji 2013, pp. 2537, 2542-4.)[3]

More broadly speaking, there are good Bayesian reasons for doubting that the coherence of the contents of testimonies is, even *ceteris paribus*, truth conducive in general. When the probability of H is updated by conditionalizing on some conjunction of evidence (E_1 &...& E_n), what ultimately affects the confirmation or disconfirmation of H is the cumulative Bayes factor of the evidence,

$$\frac{P(E_1 \& \ldots \& E_n|H)}{P(E_1 \& \ldots \& E_n|\sim H)}.$$

The unconditional coherence of the contents of a set of testimonial evidence is valuable to the confirmation of H if and only if there is some route by which it boosts the cumulative Bayes factor of the evidence in favor of H. If one prefers the *r* measure of confirmation to the *l* measure, if the prior probability of H is held constant, the *r* measure of confirmation will be higher just in case this cumulative Bayes factor of the evidence is higher.[4]

But the unconditional coherence of the contents of testimony does not *in general* boost the power of a cumulative Bayes factor for H. This is evident in a straightforward way when both H and ~H "screen off" all the E's from one another, a situation Huemer (2007, p. 340) calls full conditional independence. (See also Olsson 2005, pp. 27-28.) Where full conditional independence holds, the cumulative Bayes factor for the conjunction of the evidence is simply the product of the individual Bayes factors that the items of evidence have for H. Even if the evidence is testimonial in nature, the coherence of the contents of the testimonies has no effect of its own on the cumulative Bayes factor of the set of testimonies, because the testimonies are influencing one another only "by way of" their individual impact on H and ~H. (See Shogenji 2007 and L. McGrew 2016.)

To illustrate, suppose that Smith has been found dead, fatally shot. Suppose that two witnesses have seen the crime in sub-optimal conditions, such as at a distance or in poor lighting. Suppose that they are isolated from one another and testify independently to exactly the same content–that Jones shot Smith. Say that their individual testimonies each provide a modestly confirmatory Bayes factor of 5 to 1 that Jones indeed shot Smith, and that we can treat their testimonies as completely independent conditional on either H ("Jones shot Smith") or its negation. Now instead suppose that there is only one eyewitness of the shooting, under the same sub-optimal conditions, but that a different witness testifies independently to having overheard

[3]Throughout this discussion, I have in mind Shogenji's measure (1999) of the coherence of a set of testimonial contents, because it seems to capture well the notion of unconditional mutual positive relevance. However, as far as I can tell, my discussion in this paper does not turn essentially on the use of that measure as opposed to others thus far proposed in the literature.

[4] See Fitelson (1998, S367).

quite clearly a conversation in which Smith blackmailed Jones by threatening to reveal Jones's illegal drug activities, and Jones threatened Smith's life. Suppose that, given the background knowledge that Smith was indeed fatally shot, the testimony about the conversation also has a Bayes factor of 5 to 1 in favor of the conclusion that Jones shot Smith, and that the independence conditions are the same. In both cases, the cumulative force of the witnesses' evidence is 25 to 1 in favor of Jones's guilt. It makes no difference, either favorable or unfavorable, to the strength of the case against Jones that in the one case the witness testimonies have the same content (maximal coherence, on Bovens and Hartman's definition) and in the other case they do not. Since we have stipulated that they are independent *modulo* H and *modulo* ~H, this fact about the difference in the contents has no way to influence the cumulative Bayes factor of the testimonies.[5]

The most that can be said in a case of full conditional independence is that coherence of the contents may indicate that each of the witness testimonies has an individually confirmatory Bayes factor for the same hypothesis. We normally assume that testimonies confirm their own content. There could be odd background circumstances in which a witness's testimony that Jones shot Smith is actually evidence *against* that conclusion or in which the witness is considered to have no relevant knowledge, so his testimony has a neutral Bayes factor for its content, but these are unusual. If the testimonies each confirm their own content and if they are fully conditionally independent, the coherence of their contents stands as a proxy for the fact that the testimonies are positively relevant to the same H. But even where the unconditional coherence of the contents is *correlated* with the existence of a cumulative case for H from the testimony of conditionally independent witnesses, the coherence of the contents is not what is doing the epistemic work of confirming H. That work is being done by the force the individual testimonies have for H.[6]

If the testimonial reports are fully conditionally independent, then it is also true that there is no probabilistically separate property that one can call "coherence of the reports" that contributes to the strength of the case for H. The unconditional coherence of the reports themselves in such cases is entirely an epiphenomenon of the direction and force of the individual confirmations they provide.[7] Whether we choose to say that these are cases where "coherence" is valuable or not will depend upon how we choose to define that concept in this type of epistemic context (see section 5 on varying concepts of "coherence for H"). In any case, the epistemic evaluation of the situation will be much clearer if we focus on the reports and on the Bayes factors for H–both cumulative and individual. In the case of full conditional independence, this will lead us to note the direction of the individual Bayes factors of the reports and the fact that the cumulative Bayes factor is simply the product of those individual confirmations.

Suppose instead that full conditional independence fails and that the items of testimony are positively or negatively relevant to one another *modulo* either H or ~H or both. In that case, what Shogenji (2013, pp. 2532-3) calls the conditional coherence of the reports becomes relevant. The conditional coherence of the reports for some H_1 is

---

[5]If full conditional independence or at least independence *modulo* ~H were not specified, the greater coherence of the contents in the first scenario might actually *lessen* the strength of the case for Jones's guilt, since it would open the possibility of collusion. I will return to this point momentarily.

[6]Shogenji (2007) makes a very similar point about the individual force of information and coherence in cases of full conditional independence, though for purposes of his analysis in that article he erases the distinction between the contents of the testimonies and the testimonial reports.

[7] I thank a reviewer for suggesting that I make this point clearer.

$$\frac{P(E_1 \& \ldots \& E_n | H_1)}{P(E_1 | H_1) \times \ldots \times P(E_n | H_1)}.$$

If the reports are conditionally independent *modulo* H_1, this ratio is neutral at 1. But if conditional independence *modulo* H_1 fails, then the ratio will be either top-heavy (if the items of evidence are positively relevant to each other conditional on H_1) or bottom-heavy (if the items of evidence are negatively relevant to each other conditional on H_1).

All manner of variations in the coherence of the reports conditional on H or on ~H can influence the confirmation that the conjunction of the reports gives to H. For example, if the items of evidence (the reports) are positively relevant to each other conditional on H, this is a good thing, all else being equal, for the confirmation of H (see McGrew 2003). If the reports are positively relevant to each other conditional on ~H, this is a bad thing, all else being equal, for the confirmation of H.

That is all at the level of the reports, not of the contents. The contents of the reports *might* be positively relevant to each other in such a way that the reports themselves are unified by H (that is, are positively relevant to each other given H). That would mean that the apparent coherence of the contents was a signal of unification by H. For example, suppose that one witness testifies that Jones and Smith had a violent quarrel on the day before Jones was shot. Another witness testifies that Jones was heard muttering threats against Smith. A third witness testifies that Jones loaded a gun and went out saying that he was looking for Smith. Consider the hypothesis H, "Jones shot Smith." Given that H is true, if we hear from one witness about the threats, we may well expect more than we would otherwise expect to hear about the quarrel, the loaded gun, and the murderous intention. The hypothesis that Jones is the culprit taken together with one piece of the evidence gives us some reason to expect the other pieces of evidence. This unification which H provides is reflected in the contents as well. The reason that the *reports* are somewhat unified by H is that the *contents* are somewhat unified by H. If Jones really did shoot Smith, and if Jones and Smith really had a quarrel, there is some reason to believe that Jones was angry and threatening toward Jones and that he went out seeking him with a gun. Conversely, if Jones really did shoot Smith and really did go out seeking him with a gun, there is some reason to believe that he and Smith had had a quarrel. If we hear a report that one of these is true, we have some reason (though perhaps only weakly) to expect to hear a report that another is true as well.

A striking example in which similarity of report contents signals unification by H is the case of gravitational lensing discussed by McGrew (2003, pp. 536-565). The appearance of two spectrally identical quasar images a few arcseconds apart turned out to be the result of a massive object lensing the radiation of a single quasar and producing a double image. It is an interesting question whether we should say, strictly speaking, that the contents of the reports (the quasar images, with their locations) were coherent *unconditionally*. The prior probability of lensing was presumably quite low, so, given that the contents of the reports included the location of the images, one would scarcely think that the probability of the contents of the second report would be raised (or only to a minuscule extent) by knowing, even with certainty, the contents of the first report. Intuitively we would be inclined to say that, without some *other* confirmation of lensing, the report contents in all their detail were as near to independent as makes little difference. However, one can at least say that there were similarities in the content of the reports–namely, the identical characteristics of the images. This similarity of the contents ended up being correctly explained by L, the hypothesis of gravitational lensing of the image of the same quasar, which strongly unified the reports because it strongly unified their contents.

Individually, the reports had *no* positive relevance to L (since just one report of a particular spectral image gives us no reason to think that lensing is occurring), but jointly they confirmed L, because the conditional coherence of the reports on L was so much higher than on ~L. This case goes to show that what one can pre-theoretically think of as similarity of report contents can indicate that the items of data cumulatively support the same H even when the individual items, taken in isolation, do not confirm H.

But once full conditional independence has been abandoned, we come up against the concern that there may be unification of the reports by ~H. If the contents of three witness testimonies all seem to finger Jones as the murderer, we have to wonder if the three witnesses are colluding to produce a damning case against Jones. Given ~H and collusion, one such witness report might very well lead us to expect another. It is for this reason that Bayesian coherence theorists have often stipulated independence *modulo* ~H or similar conditions meant to rule out witness collusion. (See C. I. Lewis 1946, pp. 344, 349, Olsson 2005, pp. 27-8, 45-56, Bovens and Hartmann 2003, pp. 16, 57-61, Huemer 2007, p. 341.) Unfortunately, it is not always possible to establish that the reports are in fact independent conditional on ~H. Nor is collusion the only way in which reports may be dependent conditional on ~H, especially when H is not the conjunction of the contents, so even ruling out collusion on report contents does not rule out all dependence *modulo* ~H. For example, if one witness says that the weapon belongs to Jones, even on the assumption that Jones is innocent we may expect more strongly that another witness will say that the weapon belongs to Jones. Perhaps the weapon really does belong to Jones, but someone else used the weapon to commit the murder. (See Shogenji 2013, p. 2529.) In this way the similarity of the report contents can be explained and the reports unified by ~H. In many circumstances it would be better for the confirmation of H if the contents of the reports were *less* similar. A lesser degree of similarity in report contents can help to rule out witness collusion and other dependencies that undermine the confirmation of H. (See Shogenji 2013 and L. McGrew 2016.)

This entire discussion goes to show that, from a Bayesian perspective, the unconditional coherence of the contents of testimony cannot be *per se* a good thing for the confirmation of H. Again, what ultimately matters is the confirmation from the *set of reports* to H, and whether the coherence of the *contents* is helpful or harmful to H depends on the myriad factors that affect the confirmation from the *set of reports*. These include, *inter alia*, the force the reports have individually for H, background evidence about the probability of collusion or other ways in which the reports might be unified by ~H, and background evidence that might allow us to say that the reports are unified by H.

There is a whiff of gerrymandering about stipulating that the reports are independent conditional on ~H (or fully conditionally independent) when we are trying to decide whether coherence of the contents of reports is a positive or a negative factor for truth-conduciveness or some other desirable property. Such stipulations rule out *ab initio* the possibility that similarity of contents signals something bad for the confirmation of H.[8]

If we do not make the stipulation that the reports are independent given ~H, then at most,

[8] Erik Olsson and Stefan Schubert (2007, Schubert 2011) have proven that in a limited set of scenarios the coherence of the contents of testimony is "reliability conducive," meaning that the probability that the independent witnesses are perfectly reliable is raised by the agreement or (in Schubert's 2011 proof) the degree of overlap of their contents. These proofs are related to a point made by Rodney Holder (1998, p. 52) concerning testimonial confirmation and multiple possible false stories. The point I am stressing here is relevant to all of these–namely, that independence conditional on ~H has to be simply stipulated in order for the positive effects of coherence of *contents* to be guaranteed.

the coherence or pre-theoretic similarity of the contents can provide a heuristic clue that the reports may each be pointing positively to H (as discussed above) or that H unifies the reports more than does ~H (as in the case of gravitational lensing). But since the ultimate determinant of confirmation is the effect of the reports–individually and/or jointly–on H, the coherence of the contents may or may not be helpful to H.[9]

Olsson (2005, p. 136) puts a similar point in terms of reliability.

> ...[T]here is little hope that coherence can be truth conducive. The reason is that coherence, by its very nature, is a property on the level of report contents and is accordingly not allowed to make reference to reliability, which is a mixed-level property that concerns the relation between contents and reports. In short, the posterior is severely underdetermined by facts of coherence...

And, I would add, not only the posterior of H but the confirmation that witness reports offer to H is severely underdetermined by the coherence of the contents.

4. Coherence for H

I therefore suggest that we should switch our focus from the coherence of testimonial contents to cumulative cases from reports for an H which is not restricted to the conjunction of the contents of the reports. What makes it tempting to use phrases like "what the witnesses say hangs together" and to attribute truth-conduciveness to the coherence of testimony contents is the fact that often when the contents appear similar or are unconditionally positively relevant to each other, the testimonies themselves form a cumulative case for H. Usually this is true in virtue of the fact that they all have top-heavy individual Bayes factors for H, though the example of gravitational lensing shows that a cumulative case from similar contents can be made up of items that do not individually confirm. I suggest that we treat the fact that testimonies form a cumulative case for H as a necessary condition for their having "coherence for H" and, if we wish, include in a definition of "coherence for H" further conditions that, *ceteris paribus*, increase the strength of the cumulative case for H. (See section 5 for a suggested set of necessary and sufficient conditions for forming a cumulative case and for various suggestions of necessary and sufficient conditions for "coherence for H.")

A notion of coherence for H can be applied without any special concern for the question of whether the contents are identical to each other. Coherence for H yields no preference for situations in which report contents are identical or closer to such perfect identity. In fact, when our focus is on the confirmation of some salient H, we should not at all assume that identity of report contents is epistemically more valuable for H than varied evidence.

Suppose, for example, that you have had a battery of tests for cancer including both a biopsy and an x-ray. You may well tell your friends and family, sadly, that the results of both the

[9] Steffan Angere (2007, 2008) has suggested that coherence is useful as a heuristic guide to the truth of belief sets, which would seem to apply to the H which is the conjunction of a set of report contents. It is possible that the point I am making is part of what underlies the heuristic value Angere is discussing. However, I note that in cases of perfect agreement of report contents, one would not characterize the propositions as constituting *different* beliefs in a belief set, and even reports that are not absolutely identical but whose contents are positively relevant to each other may have some propositional contents that are identical. Since Angere's definition of a belief in a belief set associates a belief uniquely with its propositional content (2008, p. 322), it seems that Angere's discussion does not really address the issue of whether identity in propositional content in a testimonial context is advantageous.

biopsy and the x-ray "say the same thing" or "say that you have cancer," though in fact the actual test results of the biopsy refer to the appearance of cells while the results of the x-ray refer to an apparent image of a tumor in the x-ray field. The contents of the tests are not maximally coherent in the Bovens and Hartman sense; they are not identical. But, speaking informally, they "hang together" and "say the same thing" in the more important sense that they provide independent lines of support for the conclusion that you have cancer.

You would not have a stronger case for cancer if you instead had two biopsies or two x-rays of the same region, even though these would (if they both came out positive) yield more coherent test contents. Arguably, the case would be weaker. (See Shogenji 2013, pp. 2526, 2531ff.) It might be pointed out that with two biopsies or two x-rays instead, the tests are not fully conditionally independent. This is quite true. Arguably, they are dependent conditional on ~H. (See L. McGrew 2016) This fact itself underscores the misguidedness of an attempt to make coherence of contents probabilistically helpful *per se*, as discussed in section 3. If similarity of contents is, in an important set of scenarios, harmful rather than helpful to the strength of the cumulative case for the salient H, then it should not automatically be regarded as a "good" type of coherence.

Something similar is true of the contents of our sensory "reports." If you seem to smell an apple (you have a smell sensation which you remember associating with the term "apple") and at the same time have a visual sensation that you seem to remember is what an apple looks like, your smell and your sight are congruent or coherent in an important and interesting sense. They both "say the same thing." You would not have a stronger case for the presence of an apple if instead you had two *identical* (say, visual) "apple reports." The case for the presence of an apple is not weakened by the fact that the contents of the different types of sensory reports are not identical. What is useful, rather, is that you have two different reports that to some degree independently confirm the hypothesis "There is an apple in front of me." The sensory reports tell the same story, which is rendered more probable by the two of them than it would be by one alone.

5. Can we measure coherence for H?

By itself, the idea of coherence for H does not suggest just one measure for this epistemically useful property in a set of evidence. That, however, is not necessarily a bad thing, since the Bayesian coherence literature is already awash in suggested measures.

I do suggest a definition of a set's forming a cumulative case for H. I propose that a set of evidence $\{E_1,\ldots,E_n\}$ can be said to form a cumulative case for H just in case the cumulative Bayes factor of the conjunction of the set

$$\frac{P(E_1 \& \ldots \& E_n|H)}{P(E_1 \& \ldots \& E_n|\sim H)}$$

is greater than 1 and, for each member of the set of items of evidence, the probability of H given the conjunction, including that item of evidence, is greater than the probability of H given all the other members of the set of evidence, excluding that item.[10] This

[10] Again, if one prefers the *r* measure of confirmation, the first part of this suggestion can instead stipulate that the conjunction confirms H by the *r* measure. This will make no difference to which sets of evidence are said to produce a cumulative case. When the same H is in view on the same background evidence, the *r* and the *l* measure agree as to whether a given body of evidence confirms H or not, even if they do not give the same measure of that

definition means that the conjunction as a whole is positively relevant to H and also that there are no "fifth wheels" and no negatively pertinent items (in the context of the other items); each item in the body of evidence, *modulo* the others, is contributing something, however small, to the case for H.

The ratio itself–the Bayes factor for the conjunction–would then measure the strength of the cumulative case.

This definition does not entail that each item of evidence must have individual credibility, taken by itself, for H; it requires only that each item of evidence contribute something to the case in the presence of the other items of evidence in the set.[11] Hence the model described by Huemer (2007, pp. 340-1), in which individual items have no credibility for H taken alone but are positively relevant to H taken together does meet these criteria for forming a cumulative case, and the gravitational lensing case discussed above would count. This definition of a cumulative case therefore could include cases where the individual items all come together to confirm H but do not all "point to" H individually in the sense of separately confirming H. Whether or not one would want to say that these sets are "coherent for H" or not would depend upon one's definition of "coherence for H," which might include conditions besides forming a cumulative case.

One should, I suggest, make forming a cumulative case a necessary condition for "coherence for H." If one never attributes "coherence for H" to a set of evidence unless it forms a cumulative case as thus defined, one avoids counterintuitively attributing high "coherence for H" to an arbitrarily grouped set of items of evidence that has a high cumulative Bayes factor for H only in virtue of the fact that one item gives H such high confirmation that it overcomes the force of negatively relevant items in the set. One also avoids attributing high "coherence for H" to sets in which some members are irrelevant or negatively relevant to H in the presence of the others.

In trying to decide what else one might wish to include in a set of sufficient conditions of "coherence for H" and in a measure of increasing "coherence for H," I suggest that any candidate measure that is supposed to designate increasing "coherence for H" should, in the presence of plausible *ceteris paribus* conditions, increase the cumulative Bayes factor in favor of H. As potential building blocks of a definition or measure of "coherence for H," consider the following relevant ratios:

1. The individual Bayes factors the items of evidence, including testimonial reports, have for H, that is, for some individual E_i

$$\frac{P(E_i|H)}{P(E_i|\sim H)}$$

2. The cumulative Bayes factor that the entire set has for H, that is,

$$\frac{P(E_1\&\ldots\&E_n|H)}{P(E_1\&\ldots\&E_n|\sim H)}$$

confirmation. And as long as one holds constant the prior probability of H, a set of evidence that confirms H more than another set by the *r* measure will also confirm H more by the measure of the cumulative Bayes factor.

[11]One should not, however, combine the marginal positive relevances of the individual items, after all the others are taken into account, to find the force of the cumulative case. If one were, say, to take the product of all the marginal "boosts" given by each individual item to H conditional on all of the others, and if full conditional independence among the items did not hold, one might either underestimate or overestimate the force of the case as a whole.

3. The conditional coherence the items of evidence have on the assumption of H and on the assumption of ~H, respectively. That is, for some set of pieces of evidence E_1 through E_n (see McGrew 2003),

$$\frac{P(E_1 \& \ldots \& E_n|H)}{P(E_1|H)\times\ldots\times P(E_n|H)}$$ and $$\frac{P(E_1 \& \ldots \& E_n|\sim H)}{P(E_1|\sim H)\times\ldots\times P(E_n|\sim H)}$$

4. From #3 it is possible to generate a fourth measure, a comparison, by way of a ratio of ratios, of the extent to which H and ~H unify, fail to unify, or even disunify, the items of evidence.

$$\frac{P(E_1 \& \ldots \& E_n| H)/P(E_1|H)\times\ldots\times P(E_n|H)}{P(E_1 \& \ldots \& E_n|\sim H)/P(E_1|\sim H)\times\ldots\times P(E_n|\sim H)}$$

5. Gregory Wheeler et. al. (2009, 2010, 2011, 2012, 2013) and Wayne Myrvold (1996, 2003) have both suggested the following as a measure of unification of the evidence by H:

$$\frac{P(E_1 \& \ldots \& E_n|H)/P(E_1|H)\times\ldots\times P(E_n|H)}{P(E_1 \& \ldots \& E_n)/P(E_1)\times\ldots\times P(E_n).}$$

This is the ratio of the conditional coherence of the evidence on H to the unconditional coherence of the evidence.

Various candidate definitions and measures of "coherence for H" could helpfully combine the concept of a cumulative case, as defined above, with other features based on these ratios. For example,

(Coherence for H)$_1$: A set of evidence $\{E_1,\ldots,E_n\}$ has (coherence for H)$_1$ if and only if it forms a cumulative case (as defined) and, for each E_i, the individual confirmation, ratio 1, is greater than 1. Let the *ceteris paribus* conditions for this concept of coherence for H be that the prior probability of H remains the same and that ratio #4 remains the same. Stipulate that (coherence for H)$_1$ of $\{E_1',\ldots,E_n'\}$ > (coherence for H)$_1$ of $\{E_1,\ldots,E_n\}$ if and only both $\{E_1,\ldots,E_n\}$ and $\{E_1',\ldots,E_n'\}$ form a cumulative case for H, each of the individual confirmations from the items in $\{E_1',\ldots,E_n'\}$ is greater than or equal to that of the corresponding item in $\{E_1,\ldots,E_n\}$, and there is at least one item in $\{E_1',\ldots,E_n'\}$ whose individual confirmation is greater than that of the corresponding item in $\{E_1,\ldots,E_n\}$. *Ceteris paribus*, an increase in (coherence for H)$_1$ guarantees that ratio #2, the cumulative Bayes factor, rises. (Proof omitted.)

Using this definition and measure, "coherence for H" just is the fact that the set forms a cumulative case and that each item that makes up that case, by itself, points to H. "Coherence for H" increases on this definition just in the sense that some individual item or items come to point more strongly to H while none point less strongly.

(Coherence for H)$_2$: Alternatively, one might think that the term "coherence" and the concept of "increasing coherence for H" should refer to some relation *between* items of evidence that can be measured separately from and is not a mere epiphenomenon of the individual strength and

direction of evidential items, while being helpful, given plausible *ceteris paribus* conditions, to the cumulative Bayes factor.

With these considerations in mind, one can say that a set of evidence $\{E_1,...,E_n\}$ has (coherence for H)$_2$ if and only if it forms a cumulative case (as defined) and ratio #4 is greater than 1. Let the *ceteris paribus* conditions for this concept of coherence for H be that the prior probability for H remains the same and that the individual confirmations from the items of evidence (ratio 1) remain the same. This, of course, makes this the mirror image of the previous concept of "coherence for H." Stipulate that (coherence for H)$_2$ of $\{E_1',...,E_n'\}$ > (coherence for H)$_2$ of $\{E_1,...,E_n\}$ if and only if both $\{E_1,...,E_n\}$ and $\{E_1',...,E_n'\}$ form a cumulative case for H and the ratio #4 of $\{E_1',...,E_n'\}$ is greater than that of $\{E_1,...,E_n\}$. *Ceteris paribus*, an increase in (coherence for H)$_2$ guarantees that ratio #2, the cumulative Bayes factor, rises. (Proof omitted.)[12]

I note that this definition of coherence for H does not require that the items of evidence be independent conditional on ~H but only that the conditional coherence for H is greater than that for ~H, since these are the numerator and denominator, respectively, of ratio #4.

(Coherence for H)$_3$: This is a simple definition, since a cumulative case has been defined already. The conditions for a cumulative case are both necessary and sufficient for the presence of (coherence for H)$_3$. Once we have a cumulative case in hand, as defined, any factors or combination of factors that increases the cumulative Bayes factor for H (ratio 2), while the evidence continues to meet the criteria for a cumulative case, would produce an increase in (coherence for H)$_3$. If some set $\{E_1',...,E_n'\}$ meets the criteria for a cumulative case and has a higher cumulative Bayes factor for H than $\{E_1,...,E_n\}$, which also meets those criteria, then $\{E_1',...,E_n'\}$ has greater (coherence for H)$_3$ than $\{E_1,...,E_n\}$. (And if the prior probability of H is held constant, this will mean that $\{E_1',...,E_n'\}$ also confirms H more than $\{E_1,...,E_n\}$ by the *r* measure.) An increase in (coherence for H)$_3$ might be the result of an increase in ratio 1 for some item or items, ratio 4, or (if the prior probability of H is held constant) ratio 5. One of these factors could even in principle decrease, as long as the second set of evidence, in its relation to H, continued to meet the conditions for forming a cumulative case, but this decrease could be more than offset by an increase in another helpful factor so that ratio #2 increased. The idea of this definition of coherence for H and of an increase in it is that the evidence forms a cumulative case to begin with, that this *just is* what is meant by the evidence's "cohering for H," and that its forming a stronger cumulative case is therefore tantamount to its having *more* coherence for H.

These are only some applications of the principles in this paper. Other variations are possible. For example, one might abandon the attempt to define having "some" coherence for H initially beyond the formation of a cumulative case and merely define what it means for one set of evidence to have *greater* coherence for H than another along lines similar to those suggested by (coherence for H)$_1$ or (coherence for H)$_2$. The point of all of these ideas is that the items form a cumulative case for H and that anything deemed to be an increase in coherence for H increases the strength of that cumulative case, *ceteris paribu*s, as measured by ratio #2–the cumulative Bayes factor.[13]

---

[12] If one prefers ratio #5, a similar alternative for (coherence for H)$_2$ would use ratio #5 as the measure of increased coherence. Under the given *ceteris paribus* conditions, ratio #5 increases if and only if ratio #4 increases, though it is not the case that ratio #5 is greater than 1 if and only if ratio #4 is greater than 1. Proof omitted.

[13] David Glass (2007) suggests that we consider coherence between the evidence and an hypothesis H rather than coherence among items of evidence. This interesting suggestion bears some similarity to mine, but Glass does not explicitly relate his idea to the confirmation that a set of evidence provides to H, as defined by a confirmation measure such as the cumulative Bayes factor suggested here. At one point (p. 293) he implies that an

6. Looking ahead: Coherence for H and the perennial questions

Coherence for H captures the intuition that coherence is epistemically valuable. Erik Olsson pointedly asks (2013), "How can it be that we trust and rely on coherence reasoning, in everyday life and in science, when in fact coherence is not truth conducive?" The question arises naturally from the impossibility results for truth-conduciveness and the coherence of the contents of reports. But if the conjunction of the members of a set of evidence supports some hypothesis H, and if all members of that set are playing a role in that epistemic effect, then that set of evidence is indeed confirmatory of H and in that sense truth conducive for H. We rely on *cumulative cases for* hypotheses and on the factors that strengthen them.

The concept of coherence for H is especially relevant to the classic philosophical problems which motivated the Bayesian coherence project in the first place–for example, the reliability of memory and of the senses and the existence of an external world (Olsson 2005, pp. 12, 34ff, 74-6, Bovens and Hartmann 2003, pp. 26-7, Huemer 2007, pp. 342-4). Far from suggesting that we should abandon justification of these propositions "from scratch" (see Olsson 2005, pp. 173ff), the approach here opens up a possible route to the justification of belief in the external world and the general reliability of memory and the senses. The view proposed here suggests that there are resources for answering skepticism despite the failure of the Bayesian coherence quest for a separable property, coherence of the contents of reports, that is in general truth conducive.

Suppose, for example, that the H in question is something like "a generally stable external world exists, to which I have relatively reliable sensory and memorial access." In order to be justified in believing this proposition, it is not necessary to be justified in believing the conjunction of the contents of all of one's memorial or sensory appearance "reports." Indeed, it is very likely that we sometimes remember wrongly and are sometimes sensorily deceived regarding the nature of what is around us. But that need not matter too much if the goal is to have a strong case for a broadly conceived hypothesis. The bulk of our apparent sensory and memorial reports may produce a very strong cumulative case for the existence of a relatively stable external world and for the general reliability of sensory and memorial access to that world even if we are justified in *disbelieving* the gigantic conjunction of all (or even a very large number of) the items our senses and memory seem to report. In fact, one of the virtues of realism is that, fully fleshed out in the light of total evidence, it provides resources, such as an understanding of the effects of drugs, alcohol, and sleep, which permit us to explain even those sensations and memories that are chaotic or unpredictable and that do not fit together with those that immediately precede and follow them.

The question does arise whether individual, isolated experiences have *any* confirmatory value for realism. That is a contentious issue and one I will not try to resolve decisively here. As the discussion earlier of smelling, touching, and seeing an apple indicates, I am inclined to think that, like witnesses who have *some* individual credibility for what they attest, individual sensory

---

hypothesis that is "highly confirmed" by the evidence has a high *posterior* probability. Further research would be interesting to see whether Glass's suggestion of a coherence measure (p. 284) combined with his suggestion that hypotheses be ranked according to how well they cohere with the evidence (p. 287), especially if applied to sets of evidence rather than to individual items, ends up ranking hypotheses according to how strongly they are supported by the evidence as measured by the cumulative Bayes factor or for that matter by the *r* measure of confirmation or some other popular measure.

experiences do offer some support for realism. This was C.I. Lewis's idea (1946, pp. 239, 346). Even if they do not, however (as Huemer, 2007, suggests may be the case), they are capable of participating in a cumulative case for realism if they are unified by realism as measured by ratio #4–i.e., if realism unifies our sensory and memorial experience concerning some apparent object better than does its negation.[14]

Sensory and apparent-memory experiences certainly seem to have confirmatory value for realism when taken in small groups, far less than the total of all of our experiences. These groups plausibly combine, in turn, to form a still stronger cumulative case. Setting aside the question of whether an apparent visual experience of an apple confirms realism all by itself, the combination of visual, tactile, olfactory, and apparent memory experiences of an apple seems to confirm realism. So do similar sets of experiences and/or apparent memories for other objects. If this is correct, multiple subsets of our experience can "point in the same direction" in the sense of having top-heavy Bayes factors for realism. It is not even terribly far-fetched to think that groups of apparent experiences of various objects, of apparent regularities, and of the way that objects seem to interact, are unified by realism more than by its negation.[15]

A cogent argument from experience for the existence of the external world and the reliability of memory and sensation is far beyond the scope of this paper. But as we try to move forward on that project, coherence for H points away from a focus on simplified cases in which all reports say exactly the same thing, away from a focus on the conjunction of the contents of the reports, away from a focus on the unconditional mutual positive relevance of the contents of the reports, to the potential justification via a cumulative case of the types of hypotheses many philosophers have been interested in all along. Whether that project will ultimately succeed or not, its prospects can best be evaluated by asking whether the evidence fits together in the sense that many pieces of evidence help to support the same conclusion.

[14] Some "perfect" skeptical scenario that is stipulated to be empirically equivalent to realism is not the same thing as the negation of realism. Any Bayesian approach to the problem of the external world will have to deal with the issue of what one might call absolute prior probabilities–prior probabilities relative only to tautologies. These priors, perhaps based upon simplicity considerations, will be extremely important in evaluating the posterior probability of realism.

[15] The idea here is that we could apply ratio #4 to a set of such smaller sets of apparent experiences of different objects. Ultimately, all of our evidence taken together just has whatever force it has. We obviously cannot *make* it have more or less confirmatory force by considering subsets of evidence and grouping them however we please. However, as a tool for analysis, thinking of evidence in subsets and seeing the confirmatory force these have and how they can combine in a hierarchical fashion, including asking whether unification (by H more than by ~H) exists among individual items and also among subsets, seems to be a legitimate way of getting a handle on the complex issue of what force the conjunction of all evidence has for some H. Nor is this strategy relevant only to the problem of the external world. It also seems to arise when evaluating scientific theories and other hypotheses.

Works Cited

Angere, S. (2007) “The Defeasible Nature of Coherentist Justification.” *Synthese* 157:3:321-335.

Angere, S. (2008) “Coherence as a Heuristic.” *Mind* 117:465:1-26.

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

Fitelson, B. (1998) "The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity," *Philosophy of Science* 66, Supplement, "Proceedings of the 1998 Biennial Meetings of the Philosophy of Science Association," Part I, pp. S362-S378.

Huemer, M. (2007) “Weak Bayesian Coherentism.” *Synthese* 157:337-346.

Holder, R. (1998) “Hume on Miracles: Bayesian Interpretation, Multiple Testimony, and the Existence of God,” *BJPS* 49:1, 49-65.

Lewis, C. I. (1946) *An Analysis of Knowledge and Valuation*. LaSalle, IL: Open Court Publishing Co.

McGrew, T. (2003) “Confirmation, Heuristics, and Explanatory Reasoning.” *British Journal for the Philosophy of Science* 54:553-67.

McGrew, L. (2016) “Evidential Diversity and the Negation of H: A Probabilistic Account of the Value of Varied Evidence.” *Ergo* 3, No. 10.

Myrvold, W. (1996) "Bayesianism and Diverse Evidence," *Philosophy of Science* 63:661-5.

Myrvold, W. (2003) "A Bayesian Account of the Virtue of Unification," *Philosophy of Science* 70:399-423.

Olsson, E. (2005) *Against Coherence*. Clarendon Press, Oxford.

_____. (2013) "Coherentist Theories of Epistemic Justification." *The Stanford Encyclopedia of Philosophy* (Spring 2013 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2013/entries/justep-coherence/>

Olsson, E. and Schubert, S. (2007) “Reliability Conducive Measures of Confirmation.” *Synthese* 157:297-308.

Schubert, S. (2011) “Coherence and Reliability: The Case of Overlapping Testimonies.” *Erkenntnis* 74:263-275.

Shogenji, T. (1999). “Is Coherence Truth-conducive?” *Analysis,* 59: 338–345.

_____. (2007) “Why Does Coherence Appear Truth-Conducive?” Synthese 157:3, 361-72.

_____. (2013) “Coherence of the Contents and the Transmission of Probabilistic Support.” *Synthese* 190:2525-2545.

Wheeler, G. (2009) "Focused Correlation and Confirmation," *British Journal for the Philosophy of Science* 60:79-100.

Wheeler, G. (2012) "Explaining the Limits of Olsson's Impossibility Result," *The Southern Journal of Philosophy* 50:136-150.

Wheeler, G. and Scheines, R. (2010) "Causation, Association and Confirmation" Department of Philosophy. Carnegie Mellon University. Paper 418.

Wheeler, G. and Scheines, R. (2013) "Coherence and Confirmation Through Causation," *Mind* vol.122, issue 485:135-170.

Wheeler, G. and Schlosshauer, M. (2011) "Focused Correlation, Confirmation, and the Jigsaw Puzzle of Variable Evidence," *Philosophy of Science* 78:376-92.