Explanatory coherence and the impossibility of confirmation by coherence

Ted Poston[*]

January 14, 2021

Abstract

The coherence of independent reports provides a strong reason to believe that the reports are true. This plausible claim has come under attack from recent results in Bayesian epistemology. Huemer (1997), Olsson (2002, 2005), and Bovens and Hartmann (2003) prove that, under certain probabilistic conditions, coherence cannot increase the probability of the target claim. These results are taken to demonstrate that epistemic coherentism is untenable. To date no one has investigated how these Bayesian results bear on different conceptions of coherence. In this paper, I investigate these Bayesian results by using Paul Thagard's ECHO model of explanatory coherence (Thagard (2000)). Thagard's ECHO model provides a natural representation of the evidential significance of multiple independent reports. The ECHO model, in contrast to the Bayesian models, captures the power of coherence in a witness scenario. The conditions that Bayesian models found to be impossible, ECHO models naturally accommodate. This demonstrates that there are different formal tools for representing coherence. I close with a discussion of the differences between the Bayesian model and the ECHO model.

The idea that the coherence of a body of information provides a reason for that body of information has a long history. A clear early use of coherence reasoning comes from Carneades, described by Sextus Empiricus as follows:

Just as some doctors detect the genuine fever patient not from one symptom, such as an excessive pulse or a severe high temperature, but from a cluster [of symptoms], such as a high temperature as well as pulse and soreness to the touch and flushing and thirst and similar things, so too the Academic makes his judgment as to the truth by a cluster of appearances.[1]

Carneades's point is that judgment should be responsive to a mass (a cluster) of evidence, not to a single isolated report. Meinong, writing in the early 20th century, provides the following analogy of how coherence functions: "One may think of playing cards. No one of them is capable of standing by itself, but several of them, leaned against other, can serve to hold each other up."[2] Meinong's analogy suggest that a coherent body of information itself may provide reason to believe the claims even though each individual item alone does not provide a reason to believe a claim.

The intuition that coherence is a unique source of justification is widespread. Even so, a common objection is that this intuition requires an account of the nature of coherence and no such account is forthcoming. A.C. Ewing channels this common compliant writing that apart from a clear account of coherence, the theory is "reduced . . . to be mere uttering of a word, coherence, . . . rob[bing] it of almost all significance."[3] Ewing's compliant about the nature of coherence is still relevant. Even though there have been multiple formal accounts on the nature of coherence, there is no settled view about its nature.[4] But formal epistemology has made progress on the epistemology of coherence. Recent Bayesian results have shown that coherence cannot provide confirmation unless the individual evidence itself provides confirmation. These results are taken to be bad news for coherentism. Erik Olsson, for example, writes,

Coherence cannot generate credibility from scratch when applied to independent data. Some reports must have a degree of credibility that is prior to any consideration of coherence, or such agreement will fail to have any effect whatsoever on the probability of what is reported.[5]

The formal results have been framed within Bayesian models. No one has investigated how these results bare on alternative models of coherence. In this paper I

---

[1]Sextus Empiricus (2005, 37).
[2]A. Meinong *Uber Moglichkeit und Wahrscheinlichkeit* (1915), p. 465.
[3]Ewing (1934, 246)
[4]See Roche (2013) for a recent overview of formal accounts of coherence.
[5]Olsson (2005, 69).

use Paul Thagard's ECHO model of coherence to model the witness agreement scenarios centrally at issue in the Bayesian coherence literature.[6] I show that Thagard's ECHO model captures the coherentist judgment that isolated reports fail to confirm whereas multiple reports do confirm. I then discuss what is driving the difference between the ECHO models and the Bayesian models.

1 Witness Agreement and the impossibility of coherence

In this section I review the central witness agreement model and the early impossibility results.

1.1 The witness agreement model

The Bayesian coherence literature picks up on C.I. Lewis's model of coherence justification in his 1946 book *An Analysis of Knowledge and Valuation*.[7] Lewis's key observation is that coherence is best seen in the case in which multiple witnesses report that the same event has occurred. When the witnesses are independent, he claims that the agreement of the reports provides a powerful reason to accept the report. Lewis explains,

> Imagine a number of relatively unreliable witnesses who independently tell the same circumstantial story. For any one of these reports, taken singly, the extent to which it confirms what is reported may be slight. And antecedently, the probability of what is reported may also be small. But the congruence of the reports establishes a high probability of what they agree upon.[8]

Lewis's witness agreement model takes coherence to be agreement in content. Two reports cohere when they report the same event. Lewis holds that coherence is epistemically powerful when (i) the reports are independent and (ii) individually, the reports have some positive, but small, bearing on the content of the claim. Under these conditions Lewis thinks that the coherence of the reports bestows a significant probability on the claim thus supported, even if a single report has little effect.

---

[6]Thagard (2000)

[7]Lewis (1946)

[8]Lewis (1946, 346)

Laurence BonJour picks up on Lewis's witness argument model in his defense of a coherentist account of empirical justification. BonJour claims that the coherence of witness reports is powerful even if each report, on its own, has no probabilistic effect. He writes,

> What Lewis does not see, however, is that his own example shows quite convincingly that no antecedent degree of warrant or credibility is required. For as long as we are confident that the reports of the various witnesses are genuinely independent of each other, a high enough degree of coherence among them will eventually dictate the hypothesis of truth telling as the only available explanation of their agreement.[9]

BonJour posits that the positive bump in credence that any individual report provides is not essential to the power of coherence. If the reports are independent from one another then coherence alone provides a powerful reason to think the content is true. BonJour's thought is twofold: (i) rational belief is not moved by individual testimonial reports but (ii) rational belief is moved by the coherence of the individual testimonies.

1.2 Huemer's anti-coherence theorem

Whether or not BonJour is right is a crucial question for the viability of epistemic coherentism. Can coherence increase the justification of a body of claims without first requiring that those claims have some justification independent of coherence? Michael Huemer's (1997) paper attempts to answer this question by interpreting BonJour's intuition as formal constraints on a probabilistic model. In the following I briefly explain Huemer's theorem and its purported significance.[10]

Let us begin with some terminology. Let $W_{i,A}$ indicate that witness i reports A. BonJour's claim that the witness reports need no antecedent degree of credibility may be understood thusly:

No Cred: $P(A \mid W_{i,A}) = P(A)$.[11]

In contrast Lewis's model assumes that the witness reports have, at least, some small degree of credibility. That is,

---

[9]BonJour (1985, 148)

[10]Huemer (1997)

[11]There's a tacit universal quantifier ranging over witness i. This condition should be read "for any witness i the prior probability of A is unmoved by i's report that A." The other conditions below should be read similarly.

Cred: $P(A \mid W_{i,A}) > P(A)$.

In these claims, and throughout the paper, we should understand probability as rational credence.[12] **(No Cred)** specifies that a single witness report does not move rational credence. The idea is that if one lacks any relevant information about whether the witnesses report truthfully then one shouldn't change one's credence in the claim thus reported.

The next feature of BonJour's intuition is that the witnesses are genuinely independent of each other. If the witnesses are independent, then, BonJour thinks, the coherence of their reports provides a powerful reason to believe that the reports are true. This is modeled in terms of conditional independence. This is,

Conditional Independence:
(1) $P(W_{j,A} \mid W_{i,A} \wedge A) = P(W_{j,A} \mid A)$
(2) $P(W_{j,A} \mid W_{i,A} \wedge \neg A) = P(W_{j,A} \mid \neg A)$

(1) specifies that one's credence that j will report A is responsive to A alone; another witness testimony is screened off by A. (2) specifies the same thing with respect to $\neg A$. (1) and (2) are what we'd expect for witnesses who are causally independent of each other.[13]

BonJour's intuition can then be understood as saying that under the conditions of no individual credibility and conditional independence, the agreement of multiple witness reports provides powerful evidence that the reports are true. That is,

BonJour's Formal Intuition: It is possible that $P(A \mid W_{i,A} \wedge W_{j,A}) > P(A)$ even if (i) $P(A \mid W_{i,A}) = P(A)$, (ii) $P(A \mid W_{j,A}) = P(A)$, and (ii) the reports are conditionally independent.

This is intended to capture BonJour's thought that the coherence of reports can provide a reason to believe a claim even if the reports individually do not first provide such a reason.

BonJour's intuition, thus formalized, conflicts with a theorem of probability that under these conditions, the agreement of multiple reports does not change the relevant prior probability.

Huemer's theorem: $P(A \mid W_{i,A} \wedge W_{j,A}) = P(A)$ when (i) $P(A \mid W_{i,A}) = P(A)$, (ii) $P(A \mid W_{j,A}) = P(A)$, and the reports are conditionally independent.

---

[12]See Maher (2006, 2010)
[13]See Olsson (2002, 262) for a discussion on the reasons for conditional independence.

Huemer's theorem is easily proved from two consequences of (**No Cred**). First, (**No Cred**) implies that learning that A occurs does not change one's credence that witness i testifies that A. That is, the following is true:

(†) $P(A \mid W_{i,A}) = P(A) \iff P(W_{i,A} \mid A) = P(W_{i,A})$.

(†) is to be expected when one lacks any knowledge about whether witness i tracks A. If one's prior credence in A is unmoved by a report that A then one ought to think that independently learning A doesn't change one's credence that i reports A. Similarly, if learning that A doesn't change one's prior credence that i reports A then learning that i reports A doesn't change one's credence that A. Thus, (†) implies that i's report that A is like background noise with respect to A.

Second, (**No Cred**) implies one's credence that i reports A is the same conditional A or $\neg A$. That is,

(‡) $P(A \mid W_{i,A}) = P(A) \iff P(W_{i,A} \mid A) = P(W_{i,A} \mid \neg A)$.

A natural way to understand (‡) is that individual witness reports are not responsive at all to the relevant facts. Rather the relationship between the individual reports and the relevant facts is the same as the relationship between individual flips of a fair coin. Just as there is no evidential relationship between the 1st flip of a fair coin and the 2nd flip of a fair coin, there is no evidential relationship between i's report that A occurred and A's occurrence.

Given (†) and (‡), Huemer's theorem is easily proved.[14] Olsson (2005) provides a fuller discussion of the impossibility results. Olsson, among other things, extends Huemer's negative results to models that include multiple hypotheses about witness reliability. Huemer (2011) finds a different set of probabilistic conditions that is compatible with confirmation by coherence, but these conditions require abandoning both conditional independence and no individual credibility. Olsson (2017) argues that coherence should be explicated in terms of conditional independence and no individual credibility. In the following I explore Thagard's ECHO model of the powerful of multiple coherent witness reports.

2   Coherence Maximization Model

Let us examine how the witness agreement conception of coherence can be modeled within Thagard's puzzle-solving conception of coherence.[15] We start with the idea

---

[14]I leave the proof as an exercise for the reader.

[15]Thagard (2000)

that a given set of propositions may stand in either positive or negative coherence relations to each other, though some propositions may be unrelated. Coherence relations are understood in terms of the following principles:

- **E1: Symmetry** Explanatory coherence is a symmetric relation, unlike, say, conditional probability. That is, two propositions p and q cohere with each other equally.

- **E2: Explanation** (a) A hypothesis coheres with what it explains, which can either be evidence or another hypothesis; (b) hypotheses that together explain some other proposition cohere with each other; and (c) the more hypotheses it takes to explain something, the lower the degree of coherence.

- **E3: Analogy** Similar hypotheses that explain similar pieces of evidence cohere.

- **E4: Data priority** Propositions that describe the results of observations have a degree of acceptability on their own.

- **E5: Contradiction** Contradictory propositions are incoherent with each other.

- **E6: Competition** If p and q both explain a proposition, and if p and q are not explanatorily connected, then p and q are incoherent with each other. (p and q are explanatorily connected if one explains the other or if together they explain something.)

- **E7: Acceptance** The acceptability of a proposition in a system of propositions depends on its coherence with them.[16]

Principle E4, Data Priority, is a special principle akin to epistemic conservatism. The idea is that one accepts data reports and then the rest of the principles determine the evidential significance of the reports. There is a similar conception of reports in the Bayesian witness agreement model. It is assumed that one has evidence that the reports occur, but at issue is the evidential significance of the reports.[17] The question in the Bayesian models is whether coherence bears on evidential power of the reports in mass to indicate that the content of the reports is true. So applied

---

[16]Thagard (2000, 43). See also Thagard (1989); Thagard and Verbeurgt (1998)

[17]It is not often made explicit in the Bayesian witness argument models but one's credence in each witness report is 1. This tracks the standard Bayesian line that conditionalization requires that evidence is learned with certainty. The anti-coherence theorems can be proven for Jeffrey conditionalization as well.

the puzzle-solving conception of coherence, starts with a body of claims that are coherent and incoherent in various ways and attempts to determine a scoring rule that will guide which subset of these claims should be accepted and which rejected. In the following I layout a general approach to a coherence problem that explains the foundations of Thagard's ECHO model without appealing to Thagard's specific neural network algorithm.[18]

We begin with the idea that coherence and incoherence is a two-place relation between propositions. The coherence of a body of information is maximized when the positive and negative constraints are maximized. A positive explanatory constraint between two propositions is satisfied when both propositions are accepted. A negative explanatory constraint between two claims is satisfied when one is accepted and the other is rejected.

Let us examine how this works in a simple model. Consider a set of information that consists of two reports e_1 and e_2 and two hypotheses h_1 and h_2 which offer competing explanations of the evidence. We then have the following set of information: $\{e_1, e_2, h_1, h_2\}$. Hypotheses h_1 and h_2 contradict each other. h_1 explains e_1 and e_2 while h_2 explains only e_2. We then characterize a set of positive constraints and a set of negative constraints. The set of positive constraints is this: $C+ = \{(e_1, h_1), (e_2, h_1), (e_2, h_2)\}$. The set of negative constraints is this: $C- = \{(h_1, h_2)\}$. Our coherence problem is then to find a partition of E into accepted claims and rejected claims that satisfies the most constraints.

We can represent this information in terms of an undirected graph. The solid lines between nodes represent a positive explanatory constraint. The dotted line represents a negative explanatory constraint.

A simple coherence problem
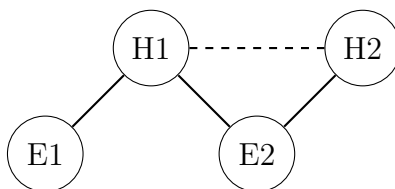
What partition of E has the highest coherence score? There are 4 propositions in E. So there are $2^4 = 16$ partitions of E into accepted and rejected items. Let us look at two partitions. First, let us consider a partition that accepts h_2 and e_2 and rejects h_1 and e_1. This is the following: $P_1 : A = \{h_2, e_2\}; R = \{h_1, e_1\}$. P_1 satisfies one positive explanatory constraint in virtue of accepting h_2 and e_2, and it satisfies

---

[18]For the specific details of his neural network algorithm see Thagard (2000, 30–34)

one negative explanatory constraint in virtue of accepting h_2 and rejecting h_1. P_1 has an explanatory coherence score of 2.

P_1 rejects e_1 and h_1 thus leaving the positive constraint between those claim unsatisfied. Consider a different partition $P_2 : A = \{h_1, e_1, e_2\}; R = \{h_2\}$. This partition satisfies two positive explanatory constraints by accepting $h_1, e_1, \& e_2$. It also satisfies the negative constraint by accepting h_1 and rejecting h_2. It thus has a higher explanatory coherence score than P_1. By inspection of the 16 possible partitions we see that P_2 has the highest coherence score. Hence, we have the most reason to accept $h_1, e_1, \& e_2$ and reject h_2.

This coherence maximization process can be done by exhaustive search among the 2^n partitions for n elements. For each partition sum all the satisfied constraints. If a partition beats all the other partitions then one has the most reason to accept its accepted elements and reject its rejected elements. The representation of constraints can be made finer by adding weights to the positive and negative constraints. If h_1 is a much better explanation of e_1 than h_2 is of the same evidence then we can represent that in terms of the relative magnitude of the numbers attached to the explanatory relation. Also, if some evidential statements have special significance this can be model in terms of the relation of the item of evidence to a unit, denoted 'special'. It then becomes a positive constraint that enters into the overall coherence score.

Thagard's ECHO model differs from this coherence maximization model only in terms of its efficiency in handling a large number of propositions and constraints. The algorithm he uses is designed to be a more efficient method of finding the partition with the highest score than an exhaustive search method. But while ECHO has the advantage of modeling a large number of constraints, it is not guaranteed to find the partition with the highest score. Furthermore, in the applications of coherence reasoning that drive our interest we can work with the simpler exhaustive search procedure.[19]

3 Coherence Maximization & Witness Agreement models

Let us consider how this coherence maximization method models a witness agreement scenario. I begin with the case of a single isolated report and then turn to multiple witness reports.

---

[19]In practice I will use various heuristics to find the partition with the highest score.

3.1 A single isolated report

The coherence maximization method begins with a body of evidence and then adds potential explanations to that body of evidence. Let us apply this in the case of a single witness report in which we know nothing about the reliability of the witness. This includes that we don't know any general background information that would provide reason to think that the witness was more or less reliable on this occasion. In this case we have an isolated witness of unknown reliability who reports that A. Consider two explanations of W's report that A. The first explanation is that the report is true and the witness is reliable on these matters. As a simplifying assumption, let us represent this conjunctive explanation as the single hypothesis that the witness is truthful. I use 'truthful' in a technical sense of 'being reliable and reporting the truth.' The second explanation is that the report is false and that the witness is misleading on these matters. Again to simplify the model, let us represent this conjunctive explanation as the single hypothesis that the witness is misleading. Then we have the following explanatory relations: (i) that W is truthful explains why W said A; (ii) that W is misleading explains why W said A; (iii) the two explanations compete with each other.

Single Witness Model

- EVIDENCE

  E1: W reports that A

- HYPOTHESES

  H1: W is truthful.
  H2: W is misleading.

- EXPLANATIONS

  X1: H1 explains E1
  X2: H2 explain E1.

- COMPETES

  C1: H1 conflicts with H2

Our set of information here is $E = \{e_1, h_1, h_2\}$. The positive constraints are: $C+ = \{(e_1, h_1), (e_1, h_2)\}$ and the negative constraints are: $C- = \{(h_1, h_2)\}$. There are 2^3 possible partitions but only 2^2 which include e_1. Since we can't accept both h_1 and h_2 we can rule out that partition and the partition in which both explanations are rejected will not satisfy any positive explanatory constraints. We are left with two partitions.

1. $P_1 : A = \{e_1, h_1\}; R = \{h_2\}$. Coherence score = 2

2. $P_2 : A = \{e_1, h_2\}; R = \{h_1\}$. Coherence score = 2

We can see that this coherence maximization model does not favor either hypothesis. This result shows that an isolated report by a witness of unknown reliability does not favor either the truth-telling hypothesis or the misleading hypothesis. It is natural to understand this result as indicating that one's credence that A is unmoved by a single isolated report. This result is precisely what would be expected given the setup that we don't know anything about the reliability of the witness.

3.2   Multiple witness reports

With a single witness report the ECHO model indicates that the report has no individual credibility. The situation changes dramatically with multiple witness reports in exactly the way that BonJour original thought anticipates. For, recall, that BonJour focuses on plausible explanations of the agreement of multiple independent witness and he reasons that eventually the agreement will be, otherwise so surprisingly, that truth-telling is the only plausible hypothesis.[20] As we'll see the evidential situation with multiple witness reports is much richer than the single witness case.

Let us describe this situation. We start with the reports of two witnesses who both report that A. Our evidence includes "W1 reports that A" and "W2 reports that A". It follows that that we have also as evidence that "W1 and W2 report the same event." BonJour's intuition included that the witnesses are independent from each other. We do not need to build in this assumption at the level of evidence; rather we add that we have as evidence that "W1 and W2 have no observed contact."

We have the following evidence set.

EVIDENCE

E1: W1 reports that A.

---

[20]BonJour (1985, 148)

E2: W2 reports that A.

E3: W1 and W2 report the same event.

E4: W1 and W2 have no observed contact.

We immediately see a difference in evidence between the case of a single witness report and the case of multiple witness reports. There are, of course, more reports. But of greater significance, there is the evidence that the reports agree and that that the witnesses do not appear to have coordinated their reports. These differences also expand the range of explanatory hypotheses.

What is the hypothesis space for multiple witness reports? As with a single witness report, we have two hypotheses corresponding to whether the witness is a truthful or misleading as understood in the technical sense given above. Also, we consider the hypothesis that the witnesses are independent from each other and the competing hypothesis that the witness are colluding. We have the following hypothesis space.

HYPOTHESES

H1: W1 is truthful.

H2: W1 is misleading.

H3: W2 is truthful.

H4: W2 is misleading.

H5: W1 and W2 are independent.

H6: W1 and W2 are colluding.

Next we specify the positive explanatory relationship. We have the following:

EXPLANATIONS

X1: H1 explains E1.

X2: H2 explains E1.

X3: H3 explains E2.

X4: H4 explains E2.

X5: H1 & H3 & H5 explain E3.

X6: H2 & H4 & H6 explain E3.

X7: H5 explains E4.

That W1 is truthful, that is in our technical sense that W1 is reliable and A is true, explains why she said A. This is similar to the above single witness model. That the witnesses are both truthful together with the fact that they are independent explains why they reported the same event. Moreover, the hypothesis that the witnesses are independent explains why we do not observe any contact between them. The hypothesis of independence thus can figure in two explanations of the evidence. Further, the hypotheses that the witnesses are misleading, that is in our technical sense that they are unreliable and that A is false, do not explain the evidence that the witnesses report the same thing. To get an explanatory connection we need to add an additional hypothesis that the witnesses are colluding, i.e., H6. But note that H6 is in tension with our evidence that the witnesses have no observed contact.

The negative explanatory relations are the following.

CONTRADICTIONS

C1: H1 conflicts with H2.

C2: H3 conflicts with H4.

C3: H5 conflicts with H6.

The model for multiple witness reports is much richer than a single witness report. We have more evidence and more hypotheses to consider. Let us then work out in detail how the ECHO model issues a verdict about which partition of the information set has the highest coherence score. Recall that the information set consists of the evidence and the potential explanations. In the multiple witness model we have this information set:

$$E = \{e_1, e_2, e_3, e_4, h_1, h_2, h_3, h_4, h_5, h_6\}$$

Given our characterization of positive and negative explanatory constraints, we have the following sets of constraints:

$$
\begin{aligned}
C+ = \{ & (e_1, h_1), (e_1, h_2), (e_2, h_3), (e_2, h_4), (e_3, h_1), (e_3, h_3), (e_3, h_5), (h_1, h_3), (h_1, h_5), \\
& (h_3, h_5), (e_3, h_2), (e_3, h_4), (e_3, h_6), (h_2, h_4), (h_2, h_6), (h_4, h_6), (e_4, h_5) \}
\end{aligned}
$$

I leave it to the reader to verify that each element of $C+$ tracks a positive explanatory constraint. We have the following negative constraints.

$$C- = \{(h_1, h_2), (h_3, h_4), (h_5, h_6)\}$$

Given E, C+, and C-, our task to consider whether there is a partition with the highest coherence score. An exhaustive search algorithm would consider each of the $2^{10} = 1024$ partitions and determine whether one has the highest coherence score. We can apply some heuristics to trim the space of partitions. One heuristic is to consider only partitions that accept all the evidence statements. This leaves us with 2^6 partitions. We can then further trim the space of partition by considering the set of negative constraints. We see, for instance, that satisfying the negative constraint requires accepting exactly one of h_5, h_6. We can then look to see if one of these hypotheses stands in more positive relations than the other. By inspection, we see that h_5 explains e_4 and h_6 doesn't. Otherwise, h_5 and h_6 stand in the same number of explanatory relations. Thus we consider the partition that accepts all the evidence, h_5, and all other claims that bear positive relations to h_5. We thus get this partition.

$$P^* : A = \{e_1, e_2, e_3, e_4, h_5, h_3, h_1\}; R = \{h_2, h_4, h_6\}$$

P^* has a coherence score of 12. By inspection, no partition has a higher coherence score. This model suggests that with an information set of E we have the most reason to accept that the witnesses are truthful and hence we should accept that their claim is true.

3.3   Discussion

The ECHO model of multiple witness reports fits BonJour's original intuition that, while an isolated report doesn't confirm the content of the report, multiple independent witness reports do confirm the report. Why does the ECHO model differ from the Bayesian model with respect to the power of coherence? To answer this question let us describe another witness agreement case in which both a Bayesian model and ECHO model are in agreement.

> **Coin-flipping witnesses:** Suppose there are a pair of witnesses, Tim and Tam, who will observe an event, E, and will report either E or not E. Tim and Tam, though, will issue their individual reports by each flipping a fair coin. If the coin lands heads then report E; otherwise report not E. Both Tim and Tam flip a coin and it lands heads for both. They both report E.

It is clear that coherence of Tim's report and Tam's report does not provide any reason to think that E is true. But this case satisfies the assumptions of no-individual credibility and conditional independence. No-individual credibility requires that $P(E \mid Tim_E) = P(E)$ and $P(E \mid Tam_E) = P(E)$, where 'Tim_E' and 'Tam_E' are that Tim reports E and Tam reports E. Conditional Independence is this claim: $P(Tim_E \mid E\&Tam_E) = P(Tim_E \mid E)$ (mutatis mudantis, for Tam's report that E). In this specific case, the Bayesian model delivers precisely the correct verdict. The agreement of Tim's report and Tam's report provides no reason to believe E.

An ECHO model of the coin-flipping witnesses does not include the hypotheses that the witnesses are truthful because the setup rules out the possibility that the reports are generated by truth-telling. Rather the relevant explanatory hypothesis for the reports are whether or not the individual coins landed heads. Accordingly, the ECHO model is the following:

EVIDENCE

E1: Tim reports that E.

E2: Tam reports that E.

E3: Tim and Tam report the same event.

HYPOTHESES

H1: Tim's coin lands heads.

H2: Tam's coin lands heads.

H3: Tim's coin lands tails.

H4: Tam's coin lands tails.

EXPLANATIONS

X1: H1 explains E1.

X2: H2 explains E2.

X3: H1 & H2 explain E3.

C1: H1 conflicts with H3.

C2: H2 conflicts with H4.

C3: H3 conflicts with E1.

C4: H4 conflicts with E2.

The reader can verify that the partition with the highest coherence score accepts H1 and H2 and rejects H3 and H4. The key is that H3 and H4 conflict with the evidence while H1 and H2 explain the evidence. Of special note is that no hypothesis in this ECHO model invokes the truth or falsity of E and hence no verdict of this model is relevant to whether E is true.[21]

The crucial difference between the coin-flipping ECHO model and the witness agreement ECHO model lies in this area. In the coin-flipping model, there are no hypotheses pertaining to witness reliability and no hypotheses that bear on the truth of E. But in the witness agreement model, there are such hypotheses. In the coin-flipping case, agreement doesn't indicate that the witnesses are reliable, but in the witness agreement case it does indicate that the witnesses are reliable. This key difference is glossed over in the Bayesian models assuming that no individual credibility holds both in the single witness case and in the multiple witness case. That is, it is a constraint on the Bayesian model that $P(A \mid W_{i,A}) = P(A)$ for each witness i. In the case where there is only one witness this is true and in the case in which there are multiple witnesses this is true. Using ECHO, though, we are led to treat the single case differently from the case involving multiple witnesses. In the single case the evidence and the explanatory hypotheses do not give us any reason to think that A is more likely to be true than not. Hence the witness's report that A doesn't provide any reason to believe A. But in the multiple witness case the evidence and explanatory hypotheses do provide us reason to think that the witnesses are reporting the truth and hence that the content is true.

The upshot of this discussion is that the assumption of no-individual credibility is too strong in the Bayesian models. The effect of coherence in the multiple witness case involves changing one's relevant conditional probabilities. Prior to learning that there are multiple witness reports in agreement, one's conditional probability that a claim is true given a single witness report is the same as the probability of the report. But after learning that independent witnesses report the same event, one

---

[21]A different ECHO model can include as hypotheses E and not E. In this case, two partitions would be tied for highest coherence score–one that includes E and one that doesn't.

is rationally moved to favor the hypothesis that the witnesses are telling the truth and in that case the assumption of no-individual credibility is false. The surprising agreement is best explained by the otherwise surprising claim that the witnesses are individually credible.[22]

4 Conclusion

The intuition that the agreement of multiple independent witnesses is evidentially powerful even if each witness lacks individual credibility is powerful. I've argued that while a Bayesian model of this case conflicts with the intuition, Thagard's ECHO model is able to capture the intuition. Moreover, reflection on the difference between the ECHO model and the Bayesian model reveals a crucial assumption in the Bayesian models that conditional probabilities relating to a witnesses credibility cannot change in response to the evidence that multiple witnesses report the same event.

References

BonJour, L. (1985). *The structure of empirical knowledge*, Harvard University Press.

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

Empiricus, S. (2005). *Sextus Empiricus: Against the Logicians*, Cambridge University Press.

Ewing, A. (1934). *Idealism: A Critical Survey*, Routledge.

Huemer, M. (1997). Probability and coherence justification, *Southern Journal of Philosophy* **35**: 463–472.

Huemer, M. (2011). Does probability theory refute coherentism?, *Journal of Philosophy* **108**: 463–72.

Lewis, C. (1946). *An Analysis of Knowledge and Valuation*, Open Court.

---

[22]This feature of the Bayesian model is related to the assumption of rigidity, that conditional probabilities are unchanged by the evidence. For a discussion of rigidity and its connection to holism see Weisberg (2009, 2015).

Maher, P. (2006). The concept of inductive probability, *Erkenntnis* **65**: 185–206.

Maher, P. (2010). Bayesian probability, *Synthese* **172**: 119–127.

Olsson, E. (2002). What is the problem of coherence and truth?, *Journal of Philosophy* **XCIX**(5): 246–272.

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

Olsson, E. (2017). Coherentist theories of epistemic justification, *in* E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, spring 2017 edn, Metaphysics Research Lab, Stanford University.

Roche, W. (2013). Coherence and probability: A probabilistic account of coherence, *in* M. Araszkiewicz and J. Savelka (eds), *Coherence: Insights from philosophy, jurisprudence and artificial intelligence*, Springer, pp. 59–91.

Thagard, P. (1989). Explanatory coherence, *Behavioral and Brain Sciences* **12**(3): 435–467.

Thagard, P. (2000). *Coherence in Thought and Action*, MIT Press.

Thagard, P. and Verbeurgt, K. (1998). Coherence as constraint satisfaction, *Cognitive Science* **22**(1): 1–24.

Weisberg, J. (2009). Commutativity or holism? a dilemma for conditionalizers, *The British Journal for the Philosophy of Science* **60**(4): 793–812.

Weisberg, J. (2015). Updating, undermining, and independence, *British Journal for the Philosophy of Science* **66**: 121–159.