

Explaining The Limits of Olsson's Impossibility Result

Gregory Wheeler

Draft of August 14, 2011

Abstract

In his groundbreaking book, "Against Coherence," Erik Olsson presents an ingenious impossibility theorem which appears to show that there is no informative relationship between probabilistic measures of coherence and higher likelihood of truth. Although Olsson's result provides an important insight into probabilistic models of epistemological coherence, the scope of his negative result is more limited than generally appreciated. The key issue is the role conditional independence conditions play within the witness testimony model Olsson uses to establish his result. Olsson maintains that his witness model yields charitable *ceteris paribus* conditions for any theory of probabilistic coherence. Not so. In fact pick, Olsson's model, like Bayesian witness models in general, selects a peculiar class of models that are in no way representative of the range of options available to coherence theorists. Recent positive results suggest that there is a way to develop a formal theory of coherence after all. What's more, although Bayesian witness models are not conducive to the truth, they are conducive to reliability.

Keywords: Bayesian epistemology, focused correlation, reliability conduciveness.

1. Introduction

One tenet of the coherence theory of justification is that more coherence entails a higher likelihood of truth, all things considered. Yet, formulating a theory of coherence to satisfy this principle is a notoriously difficult problem. Recent attempts within Bayesian epistemology to investigate the relationship between measures of probabilistic association and measures of incremental confirmation have appeared to be undermined by a series of impossibility results due to Luc Bovens and Stephan Hartmann (2003, 2005) and Erik Olsson (2005). This essay explores the limits of Olsson's result, but several of the observations made about Olsson's theorem carry over to Bovens and Hartmann's results as well.

Olsson's theorem is the strongest of the group in the sense that he views his main result to demonstrate that there is no possibility for a probabilistic measure of coherence to be 'truth-conducive' in the manner presumed by many coherence theorists. Although his result is designed to hold within a particular class of models, which we will describe in the next section, Olsson thinks this class is based on assumptions that are charitable to coherentism. For this reason, he views the negative results in this class of models to be indicative of any reasonable model of coherence:

First of all, the standard situations in which the addition of an agreeing testimony has a positive effect on the likelihood of truth are such that the reports satisfy the further conditions of being collectively independent and individually to some degree credible. And, what is more, even under such favourable circumstances, the effect of adding one more agreeing testimony on the likelihood of truth need not be very impressive, since the latter depends on the prior probability of what is being reported and also on the exact degree of credibility each witness has taken singly. So, far from guaranteeing a high likelihood of truth *by itself*, testimonial agreement can apparently do so only if the circumstances are favourable as regards independence, prior probability, and individual credibility (Olsson 2005, p.2).

Yet, Olsson's model is neither charitable to the coherence theorist nor is it representative. The result does tell us something very important about modeling coherence, however. Our aim in this essay is to highlight those points.

This paper is organized as follows. In section 2 we briefly discuss some reasons to doubt that probabilistic models of coherence are addressing the concerns of traditional coherence theories. In section 3 we review Olsson's model of witness testimony and review his impossibility theorem. In section 4 we discuss a new framework for studying the relationships between confirmation and correlation through causal structure (Wheeler and Scheines, forthcoming). This framework yields a number of strong possibility results, and those results help demarcate the limits of Olsson's theorem. Finally, in section 5, we return to the problem of

independent witness reports and introduce a reliability conducive model which offers a way around Olsson's impossibility theorem.

2. Coherentism and Bayesian Epistemology

One might challenge Bayesian epistemology in general, and Olsson's assault in particular, on the grounds that the Bayesian theory of coherence does not address the concerns of traditional coherence theories of justification. For instance, the Bayesian truth conduciveness problem, as I've discussed elsewhere (Wheeler 2009), is not the same thing as the *truth connection* problem that plagues traditional coherence theories (BonJour 2002). Bayesian truth conduciveness is solely a property of (a set of) probability functions, which may either be interpreted as degrees of belief or, less plausibly, as objective chances. But these models may not entertain both interpretations at once. Yet, the truth connection problem is concerned with the *relationship* between an agent's coherent beliefs, however one wishes to model them, and the world outside the agent's head. Since Bayesian epistemology is committed to purely probabilistic methods (Bovens and Hartmann 2003, 12), this means that a single interpretation of its probability measures will exclude either credences or objective chances even though the truth connection problem calls for a probabilistic theory to juggle both.

Second, Bayesian epistemology has treated coherence as a property of a set of propositions— it is the propositional contents of an agent's beliefs or the propositional contents of eyewitness reports which are coherent, not the beliefs or witness reports themselves. But there are doubts whether this is the central notion of coherence. For as BonJour has remarked, albeit after he renounced the coherence theory,

The fact that a belief was caused in this way rather than some other can play a crucial role in a special kind of coherentist justification. The idea is

that the justification of these perceptual or observational beliefs, rather than merely appealing to the coherence of their propositional contents with the contents of other beliefs (so that the way that the belief was produced would be justificationaly irrelevant), appeals instead to a general belief that beliefs caused in this special way (and perhaps satisfying further conditions as well) are generally true (BonJour 2002, 206-7).

BonJour is discussing how the coherence theory can explain the justification of observational beliefs, which is precisely the notion of justification that Bayesian epistemology has attempted to address. What's more, Bayesian models of witness testimony initially intended to model BonJour's theory of 'cognitively spontaneous beliefs' (BonJour 1985), which was his earlier proposal for coherentist justification of observational beliefs. Indeed, Richard Scheines and I have argued that accounting for the cause of coherence is crucial to modeling coherence (Wheeler and Schienens, forthcoming).

This said, the history of the coherence theory of justification is long on appeals to coherence justification but short on providing details for a theory *of* coherence. So, despite reservations one may have about the details of the Bayesian project, the attempt to give a probabilistic model of coherentist justification is nevertheless is the most serious effort yet to supply coherentism with a sound theory, or to uncover a good explanation for why one is not likely to materialize.

Although there are differences between Bovens and Hartmann's model for witness testimony (Bovens and Hartmann 2003) and Olsson's testimonial systems, there is an underlying structure that warrants viewing them as each of the same kind. What binds Olsson's and Bovens and Hartmann's models together, and what drives the impossibility results, is the use of conditional independence conditions to regulate the relationships between witness reports and the content of those reports (Wheeler 2009, Wheeler and Scheines, forthcoming). Seeing the impossibility results as telling us primarily about Bayesian witness models, rather than about the general relationship between (informative) measures of coherence

and truth conduciveness is an crucial insight, regardless of one's stake in coherentism.

Regarding the impossibility results, it is not the mathematics that are in dispute, but rather the philosophical license they are thought to provide. Understanding the limits of the impossibility results is an exercise worth undertaking independently of the fit between traditional coherentism and Bayesian coherentism. In the case Bovens and Hartmann's model and Olsson's Lewisian model for testimonial system, both pick out a specific class of probability models which simultaneously *overconstrain* the truth conduciveness problem, by assuming that witness testimonies are conditionally independent from one another (Wheeler 2009), and *underconstrain* the problem by assuming that probability alone will suffice (Wheeler and Scheines, 2011).

To illustrate, imagine I tell you that the world supply of scandium is running out and, after I tell you this, in an effort to provide additional evidence for this claim, I tell you that I read the news about dwindling scandium supplies in three newspapers. Now, there are several things that will influence whether my testimony will sway your opinion, but surely your position on scandium supply after my first report would not strengthen after my claim of having read the three newspaper reports if you were to learn that I read those reports from identical copies of the same edition of the same newspaper. Although all three reports are highly coherent, the coherence of the three accounts—such as it is—offers no additional reason for you to believe that the world is running out of scandium.

What goes wrong in this example? According to Olsson and Bovens and Hartmann, the problem is that the newspaper reports fail to be independent from one another. If I had succeeded in finding three news outlets which had independently sourced the story, then the coherence among the accounts would appear to give additional evidence for the claim that the scandium supply is

running low. (More on this later.) Bayesian witness models are specifically designed to rule out collusion, so Olsson, Bovens, and Hartman would handle this newspaper example by simply excluding it out of hand. What's more, although Olsson acknowledges that it is "implausible to require full independence for coherence to have the desirable effect" (Olsson 2002, p. 259), he nevertheless claims that Bayesian witness models in general, and his model in particular, specifies the most favorable circumstances.

An alternative explanation for what goes wrong does not rely upon a failure of independence, but rather pins the blame on an irrelevant *cause* of the coherence (Wheeler and Scheines 2011). The coherence of the newspaper reports is explained by facts to do with how newspapers are printed and distributed, and those facts have nothing to do with the veracity of what is printed on the page. One should not confuse a particular mechanism for constraining the possible sources of coherence—such as the conditional independence conditions of Bayesian witness models—for the general point that accounting for the source of coherence is a necessary condition for determining whether coherent evidence is a help or hindrance to epistemic justification. Indeed, if we altered the newspaper example by replacing the claim about scandium shortages with another claim about the newspaper's printing press, then those three papers could easily provide ampliative evidence after all.

3. Olsson's Result

Suppose that H is a binary variable representing a hypothesis, where ($H = 1$) codes for the proposition 'the hypothesis is true,' written H for short, and ($H = 0$) codes for the proposition 'the hypothesis is false,' written $\neg H$ for short. We will use this convention to present a generic version of the Bayesian model for witness testimony in section 4, but first consider the particulars of Olsson's model and his main result.

A *Lewis testimonial system*, $\mathbf{S} = \{(E_1, A_1), \dots, (E_n, A_n)\}$, is a set of pairs of propositions representing n witness testimonies. Each pair in the set, (E_i, A_i) , encodes an individual i 's witness report, where the proposition E_i is the report that proposition A_i is true (Olsson 2005, p. 97). A *basic Lewis scenario*, (\mathbf{S}, \mathbf{P}) , is a two-witness, single hypothesis Lewis testimonial system, $\mathbf{S} = \{(E_1, H), (E_2, H)\}$, and a set \mathbf{P} of probability distributions defined with respect to an algebra defined on the propositions encoded by this testimonial system, E_1, E_2 , and H , together with propositions regulating the witnesses' reliability, R_1, R_2, U_1, U_2 , such that P is in \mathbf{P} iff the following eleven conditions are satisfied:

1. $P(R_i) + P(U_i) = 1$, for $i = 1, 2$;
2. $0 < P(H) < 1$;
3. $P(E_1 \mid H, R_1) = 1 = P(E_2 \mid H, R_2)$;
4. $P(E_1 \mid \neg H, R_1) = 0 = P(E_2 \mid \neg H, R_2)$;
5. $P(E_1 \mid H, U_1) = P(H) = P(E_2 \mid H, U_2)$;
6. $P(E_1 \mid \neg H, U_1) = P(H) = P(E_2 \mid \neg H, U_2)$;
7. $P(R_i \mid H) = P(R_i) = P(R_i \mid \neg H)$, for $i = 1, 2$;
8. $P(U_i \mid H) = P(U_i) = P(U_i \mid \neg H)$, for $i = 1, 2$;
9. $P(E_1 \mid H) = P(E_1 \mid H, E_2)$;
10. $P(E_1 \mid \neg H) = P(E_1 \mid \neg H, E_2)$;
11. $P(R_1) = P(R_2) > 0$.

Since the reliability of each of the witnesses is assumed to be identical, $P(R_1) = P(R_2)$, in the remainder we drop the indices and simply write $P(R)$. Basic Lewis scenarios have the following properties.

- (P1) $P(H \mid E_1, E_2) = h^* = (P(H) + P(R) P(\neg H)) / (P(H) + P(R)^2 P(\neg H))$.
- (P2) Define $P(H) = h$. Then h^* as a function of h has a unique global minimum in $(0, 1)$ at $h_{\min} = P(R) / (1 + P(R))$.
- (P3) $0 < h^* < 1$.

(P4) $h^* \rightarrow 1$ as $h \rightarrow 0$.

(P5) $h_{\min} \rightarrow 0$ as $P(R) \rightarrow 0$.

(P6) $h_{\min} \rightarrow 1/2$ as $P(R) \rightarrow 1$.¹

Turning to the definition of coherence measures, Olsson's model defines coherence within the hypothesis space, rather than in the witness report (or evidence) space.² For Olsson, a coherence measure C is *informative* within a basic Lewis scenario (\mathbf{S}, \mathbf{P}) just in case there exist measures P and P' in the set \mathbf{P} such that the coherence measure $C(\mathbf{S})$ defined with respect to P is not identical to the coherence measure $C(\mathbf{S})$ defined with respect to P' , written $C_P(\mathbf{S}) \neq C_{P'}(\mathbf{S})$. Finally, C is *truth conducive* in a basic Lewis scenario (\mathbf{S}, \mathbf{P}) just in case: if $C_P(\mathbf{S}) > C_{P'}(\mathbf{S})$, then $P(\mathbf{S}) > P'(\mathbf{S})$, for all P, P' in \mathbf{P} such that $P(R) = P'(R)$.

Olsson's main result is that there are no informative coherence measures that are truth conducive, *ceteris paribus*, in a basic Lewis scenario (Olsson 2005, Appendix B).

4. Mediating Coherence and Truth Through Causal Structure

It turns out that conditions (9) and (10) of a basic Lewis scenario do the bulk of the work in Olsson's impossibility result. More generally, implicit assumptions about how witness reports and the contents of reports are related within Olsson's model can be made explicitly by a causal Bayesian network (Spirtes et. al. 2000) of a particular kind: a *single-factor common cause model*.

¹ Properties P1 and P2 are due to (Bovens et. al. 2002), and P3-P6 are from (Olsson 2005).

² Huemer (2011) refers to this as 'content determinism'. It should be noted that the Wheeler-Scheines approach does not endorse content determinism.




Figure 1: Single-Factor Common Cause Model

This model says that the truth of the hypothesis is a direct cause of each witness reporting as they do. Here the truth of H is the one and only cause for each witness (E_i) to report that H is true. All Bayesian witness models assume that witnesses are partially reliable. On Olsson’s model, when the unreliability of witnesses is not extreme, i.e., $0 < P(U) < 1$, then $P(H | E_i) > P(H)$, for $i = 1, 2$.³ This means that each witness report (E_i) is positively correlated with the content (H) of the report.

Probabilistic measures of coherence are taken to be some measure of probabilistic association. But, if we suppose that coherence is modeled by Pearson’s correlation coefficient, ρ , then the correlation between any pair of evidence variables E_i, E_j in a single factor common cause model is simply determined by the product of the correlation between H and E_i and the correlation between H and E_j .⁴ Although we are discussing coherence in the evidence space rather than the hypothesis space, the point to realize is that within common cause models the correlation of the evidence *is determined* by the relationships between the hypothesis and the evidence, once the prior probabilities of H and the E ’s are fixed (Wheeler and Scheines, forthcoming).

³ This is Observation 3.3 in (Olsson 2005). The condition $P(H | E_i) > P(H)$, for $i = 1, 2$ is labeled (A1) in Wheeler and Scheines, forthcoming, and conditional independence is (A3).

⁴ This result follows from the Danks and Glymour trek rule (2001). See proposition 4 in Wheeler and Scheines, forthcoming, which is a simplification and generalization of the Bayesian impossibility results.

There is no way for correlation to vary within this class of models. What's more, as we see below in Figure 2, a generalization of the Bayesian witness model which represents Bovens and Hartmann's model, their content variables stand in the place of our evidence variables.

So, single factor common cause models are a natural interpretation of Bayesian witness models. Furthermore, in the special case of two evidence variables and one hypothesis variable, the causal structure of the causal Bayes net entails that conditions (9) and (10) are satisfied. What's more, unless you wiggle $P(E_1)$ or $P(E_2)$, there is simply no room within the model for the correlation between E_1 and E_2 , i.e., ρ_{E_1, E_2} , to vary. Compare this observation to Olsson's remarks about the *ceteris paribus* conditions behind his witness testimony model:

We are trying to isolate one property, namely coherence, from other properties that may have an influence on the joint probability of a set of beliefs or testimonies.... Holding the degree of reliability or independence fixed is in perfect compliance with the requirement of independent variation. Fixing these aspects of scenarios does not mean imposing any constraints on what degree of coherence can be consistently attributed. (Olsson 2005, p. 118)

The point is that the class of single-factor common cause models does not allow as much flexibility in varying coherence and individual credibility as one might think.

The single-factor common cause model is but one type of causal structure, and we see that there is no role for a measure of coherence in determining how the collection of evidence supports a hypothesis in this structure. Olsson has encouraged us to think that his results tell us something about all possible models which regulate the relationship between evidence and hypothesis. But, this is not true. For instance, recall the causal structure in our scandium shortage example from above, where the coherence of three copies of the same newspaper has nothing to do with whether in fact there is a scandium shortage.

Alternatively, if an additional report were collected from a different newspaper, we might think that this alternative report should provide more confirmation for the hypothesis than another copy of the same newspaper, all else equal, and we have results which confirm that this conjecture holds (Wheeler and Scheines, forthcoming). But here again it is the non-parametric causal structure regulating the relationships between the evidence and the hypothesis rather than probabilistic facts about measures of coherence.

Interestingly, Bayesian witness models try to account for the role of causal structure by simply looking at one class of models, and one might think of this as the extreme case where the agreement between two witnesses can best be explained by the truth of the hypothesis they are reporting on because there is no interaction between the two witnesses. While that view accords with intuitions, the step from a small interaction between witnesses to no interaction between witnesses marks a world of difference.

The main point is this: one cannot isolate coherence from the causal structure regulating the relationships among the variables designated to cohere (our E 's) and the variable(s) designated for a boost (our H). On our theory, we measure the coherence (correlation) of evidence in order to determine what, if any, boost this evidence provides to a particular hypothesis. This structure leaves open how to interpret those variables, just as it leaves open what interpretation of probability to adopt and how to view the causal structure. For example, a traditional coherence theorist, concerned to preserve the internalist nature of coherence justification, may opt for a subjective interpretation of the model, pairing a subjective interpretation of probability with an interpretation of causal structure as that agent's judgment about which causal relationships constrain his degrees of belief. On the other hand, a mathematical psychologist interested in fitting models for human decision making to experimental data might start with correlation tables generated from experimental subjects and adopt an objective

interpretation of the model. I don't wish to dismiss the importance of working through the issues that arise from how to interpret the model, but I do wish to stress that these are secondary concerns. At this early stage, we first need to understand the fundamentals of how measures of probabilistic association, causal structure, and incremental confirmation can fit together before we dive into these details. Arguably, it was by leaping straight into the philosophical details of modeling Bonjour's theory of cognitively spontaneous beliefs which has led to mistaking a narrow result about a class of models for an impossibility result about the very idea of using probability to model coherentist justification.

As an aside, by adding causal structure to the model we are taking a step in the direction of the traditional truth connection problem. Since the ampliative effect of coherence depends on how the evidence and hypothesis variables are causally related to one another, including causal structure in the model may be thought of as a proxy for beliefs about how the environment is arranged. There is a truth to the matter of how things are causally related, and an agent may harbor causal beliefs which are mistaken. Analogously, Herb Simon's theory of bounded rationality (1956) maintains that structural constraints imposed by the environment should be factored into our notions of human rationality and norms. Although mathematical psychologists have typically focused on the cognitive limitations of Simon's proposal, such as limited memory and limited processing time, introducing a role for causal structure, while not sufficient to represent all of Simon's ideas about ecological rationality, is nevertheless an advance from the common use of linear models.

Returning to Bayesian witness models, Figure 2 represents another witness model where each R_i is a binary fact variable, Rep_i is a binary report variable, and H a (latent) binary hypothesis variable.




Figure 2: Generic Common Cause Bayesian Witness Model

In (Wheeler and Scheines, forthcoming), we argue that Figure 2 is the most plausible interpretation of Bovens and Hartmann’s witness testimony model (2002). Bovens and Hartmann, like Olsson, define coherence as a property of a set of propositions, which in this model is determined by assigning either $R_i = 0$ or $R_i = 1$ (for $i = 1, \dots, n$), to the set of fact variables $\{R_1, \dots, R_n\}$. The natural interpretation of Bovens and Hartmann’s model is that the facts are the facts not merely because a witness says so, but because the truth of the matter (H) is a direct cause of those facts. Even so, for the purposes of the discussion above about the properties of single-factor common-cause models, Figure 2 can be simplified to Figure 1. For a full discussion, see (Wheeler and Scheines, forthcoming).

Causal structure is not the whole story, of course. Consider Figure 3.




Figure 3: An Imperfect Witness Model

Instead of causal structure regulating the relationship between evidence and hypothesis, here we may ask whether the correlation (coherence) between evidence variables (evidence propositions) is primarily due to the truth of the hypothesis, or primarily due to a causal relationship between E_1 and E_2 . In models of this kind, pace Olsson, there is an informative probabilistic measure called *focused correlation* (Myrvold 1996, Wheeler 2009), which robustly tracks incremental confirmation (Wheeler and Scheines, forthcoming, Schlosshauer and Wheeler, 2011). The focused correlation of a set of evidence $\mathbf{E} = \{E_1, \dots, E_n\}$ with respect to a hypothesis H is the ratio of the association (coherence) of the evidence conditional on H to the association (coherence) of the evidence *simpliciter*, which can be expressed generally as:

$$For_H(E_1, \dots, E_n) =_{df} \frac{S(E_1, \dots, E_n | H)}{S(E_1, \dots, E_n)} = \frac{P(H | E_1, \dots, E_n) P(H)^{n-1}}{P(H | E_1) \dots P(H | E_n)}.$$

What focused correlation measures is the difference between the association (coherence) in the evidence alone and the association (coherence) in the evidence when a hypothesis is true. When that ratio is greater than one, i.e., when there is greater evidential coherence given the hypothesis than evidential coherence alone, and evidence is positive and either of equal or of varying strength (within certain bounds), then we know that the occurrence of the hypothesis is a contributing cause to the co-occurrence of the evidence. We say in this case that focused correlation is *inflationary*. Focused correlation is deflationary when that ratio is less than one. If evidence is negative and of equal or varying strength, then in such cases the difference in this ratio indicates that the occurrence of the hypothesis makes the co-occurrence of the evidence less likely.

Focused correlation is applicable within imperfect witness models, which are models in which (9) and (10) of Olsson's testimonial systems do not hold.⁵ Moreover, focused correlation can also be used to compare evidence sets of size two such that an evidence set E gives more confirmation to H than an evidence set E' iff E has more focused correlation with respect to H than E' , all things equal.⁶

5. Independent Evidence and Reliability Conduciveness.

The surprise is that modeling coherence within Bayesian witness models coherence leads to ruin, but once you drop conditional independence a world of possibilities opens. Although intuition suggests that Figure 1 is the ideal case for coherentism and Figure 2 less so, the truth is exactly the opposite. The bounty of possibilities for a purely probabilistic account of coherence depends on an imperfect witness model. And while things improve as the correlation between the evidence goes to zero, at zero we switch from an imperfect witness model to a Bayesian witness model satisfying conditional independence and the possible suddenly becomes the impossible.

⁵ Because correlation is essentially a binary relation (or, in conditional form, a three-place relation), focused correlation in its general form is not a correlation measure but instead is a measure of deviation from independence. For primarily this reason, our basic possibility results hold for comparisons of evidence sets of size 2. There are ways to extend our techniques to cover larger evidence sets, but there is no automatic way to do this from probabilistic information alone. This is the main reason why we have focused on Olsson in this essay, rather than address Bovens and Hartmann's impossibility results. For while our general points about single-factor common cause models and the constraints this class of models imposes carries over to Bovens and Hartmann project as well, we need to have a policy for handling larger evidence sets before fully addressing their results.

⁶ The comparison assumes, like Olsson, positive evidence and that all evidence is of equal relevance (Wheeler and Scheines, 2011). However, the second assumption can be weakened for a more general result, namely that evidence may vary in strength within certain bounds (Schlosshauer and Wheeler, 2011).

This sudden switch is disturbing. Although we have avoided the sweeping conclusion of Olsson's impossibility result, the price seems to be a sharp discontinuity between imperfect witness models and Bayesian witness models. Call this an *irascibility result*.

However, all is not lost. Return to the basic witness model involving a single hypothesis and two evidence reports. The driving idea is that each evidence report on its own lends credence to the hypothesis, but integrating both sources of evidence offers an additional epistemic advantage. If the goal is model Bonjour's original proposal for handling sensory perceptions and the observational beliefs based on those perceptions and to explain how integrated observational evidence is superior to pieces of evidence standing on their own, then perhaps the quantity to focus on is not the proportion of true reports *per se*, but the reliability of those perceptual inputs in forming accurate observational beliefs. Here we are keeping the underlying structure of Olsson's basic Lewis scenario and its signature characteristic, evidential independence. We are also keeping in sight Bonjour's notion of cognitively spontaneous beliefs, since this proposal is what has motivated Bayesian epistemologists to explain how multiple independent information sources can increase the likelihood that an observational belief is true. Finally, we are also addressing the traditional coherence theories truth connection problem of explaining the connection between properties of this cognitive model and the content of the observational report. However, we are changing one thing. In place of Olsson's truth conduciveness property, I will substitute something I call *reliability conduciveness*.

To support this substitution, we will need a richer model. So, instead of binary random variables, our model will be built with continuous random variables. We define the probability that a continuous random variable Y takes a real value y in

the interval $[a,b]$ in the standard way, by integrating the probability density function over the interval $[a,b]$:

$$P(a < Y < b) = \int_a^b f(y)dy. \quad (1)$$

In this model the goal of the evidential report $E_1 = e_1$ is to accurately represent an environmental parameter θ , which the brain does not have direct access to. So, we view each evidence variable to be ‘truth plus noise’, i.e., $e_1 = \theta + \varepsilon_1$, where ε_1 is normally distributed error with a mean of zero and variance of σ_1^2 , and similarly for $e_2 = \theta + \varepsilon_2$. Thus, $P_{E_1}(e_1) = x$ expresses that the probability that the truth, θ , is within the interval $e_1 = (\mu, \sigma_1^2)$ is x . So, the proposal is to replace truth conduciveness, which is concerned with boosting x , with *reliability conduciveness*, which aims to shrink the variance of error in an evidential report. An observational belief or hypothesis, \hat{h} , is understood then to be an estimate of θ that arises from integrating all evidence reports, e_1, \dots, e_n .

Let us focus on a simple model involving two reports, e_1, e_2 , the truth parameter θ , and our estimate of the truth, \hat{h} . One benefit from assuming that individual evidential reports are independent is that their associated noise terms are assumed to be independent, too. But given that the noise of the evidence reports is independent and Gaussian, the estimate \hat{h} of θ which has the lowest variance is the *maximum likelihood estimate* (MLE) (Ernst and Bühlhoff, 2004). The optimal estimate of θ , our best observational belief \hat{h} based on the evidence, is a weighted average of the two evidential reports, e_1 and e_2 :

$$\hat{h} = w_1e_1 + w_2e_2, \quad (2)$$

where $w_1+w_2 = 1$, and

$$w_1 = \frac{1/\sigma_1^2}{1/\sigma_1^2 + 1/\sigma_2^2}, \text{ and } w_2 = \frac{1/\sigma_2^2}{1/\sigma_1^2 + 1/\sigma_2^2}.^7 \quad (3)$$

The inverse variance of each evidence report in equation 3 is defined as the *reliability* r of that report,

$$r_i = \frac{1}{\sigma_i^2}, \text{ for } i = 1, 2. \quad (4)$$

It follows that the reliability of the integrated estimate \hat{h} is the sum of the reliabilities of the individual evidence reports r_1 and r_2 .

In short, there is always a benefit from combining individual, independent evidence: it increases reliability. Constructing \hat{h} by the MLE method will always yield a more reliable estimate of the truth, θ , than either evidence report alone.

References

- H. B. Barlow, T. P. Kaushal, M. Hawken, and A. J. Parker, Human contrast discrimination and the threshold of cortical neurons, *Journal of Optical Society of America A – Optics, Image Science, and Vision*, 4(12), 1987, 2366-71.
- L. BonJour, *The Structure of Empirical Knowledge*. Cambridge, MA, Harvard University Press, 1985.
- L. BonJour, *Epistemology*. Rowman and Littlefield, Oxford, 2002.

⁷ This approach to constructing the optimal estimator \hat{h} is version of the Kalman filter (Kalman 1960).

- L. Bovens, B. Fitelson, S. Hartmann, and J. Snyder, Too Odd (not) to Be True, *The British Journal for the Philosophy of Science*, 53, 2002, 539-63.
- L. Bovens, and S. Hartmann, *Bayesian Epistemology*. Oxford University Press, 2003a
- L. Bovens, and S. Hartmann, Solving the Riddle of Coherence, *Mind*, 112, 2003b, 601-633.
- L. Bovens, and S. Hartmann, An impossibility result for coherence rankings. *Philosophical Studies*, 128, 2006, 77-91.
- C. Cross, Coherence and truth conducive justification, *Analysis*, 59, 1999, 186-93.
- D. Danks and C. Glymour, Linearity properties of Bayes nets with binary variables, in *Uncertainty in Artificial Intelligence: Proceedings of the 17th conference (UAI 2001)*, Breese and Koller (eds.), 2001, 98-104.
- M. O. Ernst and H. H. Bülhoff, Merging the senses into a robust precept, *Trends in Cognitive Science*, 8(3), 2004, 162-169.
- M. Huemer, Does probability theory refute coherentism? *Journal of Philosophy*, 108, 2011, 35-54.
- C. I. Lewis, *An Analysis of Knowledge and Valuation*. Open Court, La Salle, 1946.
- R. E. Kalman, A new approach to linear filtering and prediction problems, *Journal of Basic Engineering, Series D*, 82, 1960, 35-45.
- W. Myrvold, Bayesianism and Diverse Evidence: A reply to Andrew Wayne, *Philosophy of Science*, 63, 1996, 661-5.
- E. Olsson, What is the problem of coherence and truth? *Journal of Philosophy*, 94, 2002, 246-72.
- E. Olsson, *Against Coherence: Truth, Probability and Justification*. Oxford University Press, Oxford, 2005
- M. Schlosshauer and G. Wheeler, Focused correlation, confirmation, and the jigsaw puzzle of variable evidence, *Philosophy of Science*, 78(3), 2011, 376-92.
- H. Simon, Rational choice and the structure of the environment, *Psychological Review*, 63, 1956, 129-38.
- P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search*. 2nd. edition. MIT Press, Cambridge, MA, 2000.
- G. Wheeler, Focused correlation and confirmation. *The British Journal for the Philosophy of Science*, 60(1), 2009, 79-100.

- G. Wheeler and R. Scheines, "Causation, Association, and Confirmation," in *Explanation, Prediction, and Confirmation. New Trends and Old Ones Reconsidered*, D. Dieks, W.J. Gonzalez, S. Hartmann, T. Uebel, M. Weber (eds.), The Philosophy of Science in a European Perspective Series, Dordrecht: Springer, 2011, 37-51.
- G. Wheeler and R. Scheines, "Coherence and Confirmation Through Causation," manuscript under review.