Formal Epistemology

Gregory Wheeler ©
New University of Lisbon
Draft of July 31, 2011

Dedicated in memory of Horacio Arló-Costa

Narrowly construed, formal epistemology is a methodological approach to traditional analytic epistemology. According to this view, the aim of formal epistemology is to harness the power of formal methods to bring rigor and clarity to philosophical analysis.

Yet, in broader terms, formal epistemology is not merely a methodological tool for epistemologists, but a discipline in its own right. On this programmatic view, formal epistemology is an interdisciplinary research program that covers work by philosophers, mathematicians, computer scientists, statisticians, psychologists, operations researchers, and economists who aim to give mathematical and sometimes computational representations of, along with sound strategies for reasoning about, knowledge, belief, judgment and decision making.

This essay presents a two-pronged argument for formal epistemology. The first part addresses the general question of why anyone should bother with formal methods by illustrating, through a historical example, the role that formal models can play in inquiry. The second part describes two specific examples of recent work within formal epistemology, one that addresses a longstanding issue within traditional epistemology—namely, what to make of coherentist justification—and another addressing a fallacy of probabilistic reasoning which has implications across a wide range of disciplines, and thereby making a case for a broader, programmatic view. Finally, we close with a methodological proposal for epistemology, one that incorporates formal, experimental, and traditional approaches into one program.

Why be formal?

When you fiddle with the controls of your computer to select a color, either by typing a triplet of numbers to specify how much red, green, and blue to include, or by selecting a point within a color wheel, you are making use of a model that dates back to 1802. It was that year, in a London lecture hall, that Thomas Young first speculated that human color perception involves three different receptors in the eye, and a half-century later Hermann Helmholtz, that nineteenth century colossus, reckoned that each receptor was sensitive to three distinct frequencies of light, corresponding roughly to our perceptions of red, green, and blue (Helmholtz 1860). What's more, Helmholtz proposed a mathematical model whereby a good portion of the visible color spectrum can be represented by an additive mixture of those three basic frequencies. Although it took another century to confirm the physiological claims of the Young-Helmholtz's theory of color vision, in that span of time the RBG additive color model spurred the development of color photography, halftone color printing, and color television, among other things.

The story of the RBG color model offers an allegory for how to think about formal epistemology. For the story illustrates the impact a formal model can have on inquiry, and it also highlights how to assess the merits of a formal model. But one may reasonable ask whether the analogy holds for epistemology. There is, after all, plenty of epistemology that is not informed by formal methods, and plenty of epistemologists seem committed to keeping it that way. Giants have made their mark with little more than commonsense and a smidgeon of logic. Why, one might ask, bother being formal?

The best short answer to this question is one given years ago by Rich Thomason.[1] Thomason, commenting on philosophers who view formal methods as a distraction to real philosophical advancement, observed that the only real advantage that we have over the great philosophers of the past are the new methods that we have at our disposal. First-order logic. Calculus. The number zero. It is hard to imagine improving on Aristotle without resorting to methods that were simply unavailable to him. Knowing just this much about history, a better question is this: why limit your options?

To begin a longer answer, return to the Young-Helmholtz theory and notice three important stages in its development. First, there was Young's idea that people with normal color perception rely on three receptors in their eyes. This was a great idea, but there was neither a model nor much empirical evidence for it, so the idea languished for 50 years. At that time common wisdom held that those three receptors would surely be sensitive to red, *yellow*, and blue, since Newton's color wheel designated those three as the primary colors, and painters for centuries had mixed their paints from red, yellow, and blue. But intuition and tradition notwithstanding, Young's idea went nowhere until Helmholtz came along. His ingenious contribution was to run experiments in which subjects were instructed to match the color of a swatch to a color of their choice, which they selected by mixing three wavelengths of light. He also observed that the subjects could not make a match with only two of those three basic light sources. Helmholtz's experiments contributed important empirical evidence for the trichromatic account, mainly by offering grounds for replacing yellow with green in his account, but it was his introduction of a *mathematical model* to represent visible color as an additive

[1] Krister Segerberg recounts Thomason's remarks in (Segerberg 2005, p. 166).

mixture of three basic colors that really set things in motion. Helmholtz's mathematical model marked the beginning of the second stage of the theory, which spurred a number of innovations. James Maxwell introduced the first trichromatic color photograph within a decade of Helmholtz's work. Halftone printing presses came along soon after. Color television and color motion pictures followed several decades later. And yet, all of these developments occurred well before the last stage, when physiological mechanism that underpins the Young-Helmholz trichromatic theory of color vision was finally confirmed.

What is striking about this schematic history of the Young-Helmholz theory of color vision is not merely the outsized impact of the RGB color model had on the theory and on other developments far afield, but the timing and sequence of events. Let us turn to drawing out two of those features and consider their bearing on epistemology.

Experimental evidence is important but far from sufficient.

Some philosophers have attacked a practice within traditional analytic epistemology of justifying normative claims by appeals to intuition—a methodology that Weinberg, Nichols, and Stich have dubbed 'Intuition Driven Romanticism (Weinberg *et. al*. 2001). For these critics, the first part of the Young-Hemlhotz history will resonant, for Young's insight was stalled because people at that time were mislead by their intuitions about the three basic 'colors' necessary for color vision. It was Helmholtz's experiments that set the program on track by identifying the correct set of colors to base the theory on.

Even so, the second half of this history complicates matters for friends of experimental philosophy. For although Helmholtz's experiments were crucial for pointing to the right set of basic colors, the experimental evidence was far from conclusive for establishing the central physiological thesis about the photoreceptors in the human eye. The role these initial experiments played was to shift the focus from an unworkable RYB model to a *plausible* RGB additive color model, not to nail down the RGB model before any other progress could be made.

But look closely at Weinberg, Nichols, and Stich's attack. What they are up to, oddly enough, is refashioning a longstanding criticism of experimental psychology over *convenience sampling* (Carlson 1971, Sears 1986, Henrich *et. al*. 2010) into an attack on traditional analytic epistemology. The problem with convenience sampling in social psychological research in particular is the problem of basing general conclusions about human nature on results garnered from studies of undergraduate psychology students, who are overwhelmingly WEIRD people (Henrich *et. al*. 2010): members of Western, Educated, Industrialized, Rich, and Democratic societies. In a comparative review of the literature, Henrich, Heine, and Norenzayan found that WEIRD people are often the *least* representative populations in studies of visual perception, fairness, cooperation, spatial reasoning, categorization and inferential induction, moral reasoning, reasoning styles, self-concepts, and the heritability of IQ. They remark that

> there are no obvious *a priori* grounds for claiming that a particular behavioral phenomenon is universal based on sampling from a single subpopulation. Overall, these empirical patterns suggest that we need to be less cavalier in addressing questions of *human* nature on the basis of data drawn from this particularly thin, and rather unusual, slice of humanity (Henrich et. al. 2010, p. 61).

So, while the knock against traditional analytic epistemology is that there is experimental evidence to suggest that there are differences in epistemic intuitions in different groups and these differences undermine claims that epistemic intuitions are universal (Weinberg et. al. 2001), the knock against experimental psychology is that there are differences between the subpopulation used for the lion's share of studies and the general population, and these differences undermine 'species-level' generalizations (Henrich et. al. 2010).

Weinberg, Nichols and Stich would have us believe that different epistemic intuitions about Gettier cases (Gettier 1963) they have observed among students from East Asian or the Indian sub-continent, on the one hand, and WEIRD students on the other, undermines epistemic intuitions. (NB: It is unclear what they make of 'dual process' theories within cognitive and social psychology,[2] which distinguish between intuitive and deliberative judgments and are susceptible to an analogous WEIRD critique of Henrich *et. al*. Criticism of the dual-process theories within psychology does not involve explaining away intuitive judgments; the dispute is instead over how to provide a psychological explanation for intuitive judgments (Kruglanski and Gigerenzer 2011).)

But the problem with this line of attack on epistemic intuition is that it cloaks a weakness that experimental philosophy shares with traditional analytic epistemology, namely how to warrant a general claim—which is a necessary condition for carrying out a positive program. None of the critics would deny there are universal epistemic norms or deny there are general facts about human nature. Rather, what each critique challenges is the shared practice of making cavalier claims of generality from non-representative samples of the population. In philosophy, it is the intuitions of self-selected, highly trained, high socioeconomic status professors of philosophy that are called into question. In psychology it is US college sophomores. Even though experimental philosophy has a point in its critique of traditional epistemology, the disappointing fact remains that experimental philosophy is susceptible to the very same criticism for the general claims it would prefer we adopt instead.

[2] For overview of dual process theories, see Kruglanski and Orehek (2007), Evans (2008) and Kruglanski and Gigerenzer (2011).

The introduction of formal models can help on this front, but before considering how, let us turn to another feature of the Young-Helmholtz history.

Counterexamples are not always decisive.

Traditional epistemology, in addition to its commitment to epistemic intuitions, also tends to rely solely on counterexamples to assess theories, which means that a miss is as good as a mile. But it does not have to be this way.

Look again at the Young-Helmhotz history, for the RGB additive color model is by no means comprehensive: there are visible colors which cannot be represented within the model. Yet this limitation has never been seriously considered to be a 'counterexample' to the RGB model, nor has it been viewed as a threat to the physiological theory of vision that the model was originally designed to serve. Rather, the exceptions to the theory were weighed against the power of the model, and a recognition that the empirical basis of the theory—which remained unconfirmed through the last half of the 19th century and the first half of the 20th—was likely to be in the neighborhood of the truth. A successful formal model can make it sensible to view one theory as standing closer to the truth than another.

If experimental philosophers' complaint about epistemic intuition boils down to a beef about philosophers putting their intuitions about empirical claims above any experimental evidence to the contrary, and having the nerve to offer an *a priori* argument to justify continuing to do so, then the complaint about unchecked counterexamples boils down to a beef against constructing theories to resist objections—no matter how contrived—above all other considerations.

The problem with aspiring to counterexample-proof philosophy without taking into account either formal or empirical constraints is that the exercise can quickly devolve into a battle of wits rather than a battle of ideas. What's more, the problem is compounded by pseudo-formal philosophy—the unfortunate practice of using formal logic informally—because this encourages philosophers to describe rather than define the fundamental operations of their theories. Memories are 'accessed in the right way'; justified beliefs are 'based' on one's 'evidence'; coherent beliefs 'hang together'. But, like a bump in a rug carefully pushed from one corner of a crowded room to another, this reliance on pseudo-formalisms to avoid any and all counterexamples often means that the hard, unsolved philosophical problems are artfully avoided rather than addressed head on. At its worst, rampant counterexample avoidance threatens to turn philosophy into little more than a performance art.

But, one way to arrest this slide is by constraining epistemological theories by empirical considerations and formal models. For if you replace those fudged terms with a formal model, and hem in imagination by empirical constraints, then if a theory is successful in handling a range of cases, that hard won success will be weighed against the theory's failings. In other words, if we set aspirations for epistemology higher than conceptual analysis, that will open more room to judge success and failure than the all-or-nothing stakes of counterexample avoidance. That is one lesson of the RGB model. Prior to its introduction, people only had a vague idea of how to 'combine' colors and it was not recognized that there was a crucial difference between a *subtractive* color model, which models how paints mix to create new colors, and an *additive* color model, which is appropriate for modeling color perception. That insight far outweighed the limitations of the model, and it is the main reason that exceptions to the theory did not undermine it.

What formal models can do.

So far we have been discussing the merits of formal methods by drawing analogies to a historical example, selected purely for illustrative purposes, and by making some critical comparisons with traditional analytic epistemology and experimental philosophy's take on epistemic intuitions. There is one last piece of stage setting, which returns to a distinction we introduced at the very beginning between viewing formal epistemology as a methodological approach within analytic epistemology and viewing formal epistemology as an interdisciplinary research program, which I've discussed elsewhere—in terms that may lead to confusion—as a type of methodological naturalism (Wheeler and Pereira 2005). Here I want to simply point out the advantages to epistemologists from embracing this broader, programmatic view—whatever you would prefer to call it.

Sometimes a formal technique is used in several fields to model a family of problems, and in these cases there is often an opportunity for a formal epistemologist to build a rich repertoire of similarly structured problems together with a library of techniques for how to work with them. Probabilistic methods are an excellent example (Haenni *et. al*. 2011). But in addition to fixing the method and varying the problem, one may also focus on a single problem, which can appear in different guises in various disciplines, and vary the methods. An advantage of viewing the same problem through the lens of different models is that we can often begin to identify which features of the problem are enduring and which are artifacts of our particular methods or background assumptions. Because abstraction is a license for us to ignore information, looking at several approaches to modeling a problem can give you insight into what is important to keep and what is noise to ignore. Moreover, discovering robust features of a problem, when it happens, can reshape your intuitions. In this way formal epistemology can be used to train philosophical intuitions rather than simply serve as a tool for rigorous exposition of prior intuitions. Here then is a partial reply to the problem that besets traditional epistemology, with its reliance on epistemic intuitions, and experimental philosophy, with similar limits from relying on too many WEIRD people. The entire load of a theory need not rest entirely on the grounds for its claims if the theory includes a reasonable model that gives us new abilities to predict and explain. A good model will be put to use.

While this essay presents a case for formal epistemology, it should be clear that this is hardly a manifesto. There is a place for experimental work in philosophy, and there is a place for intuitions, too. Moreover, formal methods are not the only route to precision, and it may well be that understanding the most important human undertakings—love, friendship, political compromise—is hampered rather than helped by placing too high a stock in precision. Live long enough and you'll discover that not everything yields to hard thought.

Coherence and Dilation

So far we have discussed formal epistemology from a bird's-eye point of view. We agreed with the spirit of experimental philosophy's critique of traditional analytic philosophy, namely that unaided intuition is not likely to take us very far, but we also found experimental philosophy falling short. Throughout this discussion we used a historical example to illustrate how one might think of experimental methods and formal methods fitting together.

In this section, we shift focus to consider two recent examples of work within formal epistemology. The first example reports a breakthrough in figuring out some of the fundamentals for a theory of coherence, which is a longstanding open problem for the coherence theory of justification. Here is a classic example of formal methods being brought to bear on a problem within traditional analytic epistemology. The second example examines principles of sound probabilistic reasoning and the role that independence assumptions play. This is an example of formal epistemology pursued as a stand alone discipline, for the ramifications from this example affect the application of Bayesian methods within philosophy and far afield.

Toward a theory of coherence

In 1985, Laurence BonJour provided some structure to the coherence theory of justification,[3] and his postulates for coherentism (1985, pp. 95-9) describe a role for probability along the lines of C. I. Lewis's probabilistic model of 'congruent', self-justifying memories (Lewis 1946). Since then several authors working within the framework of Bayesian epistemology have explored the prospects of developing a probabilistic model of coherence along this basic Lewis-BonJour outline.[4]

Much of the work in Bayesian epistemology concerns coherence among a set of propositions and whether a probabilistic measure of coherence can be adduced

[3] He later despaired of meeting those demands and quit the theory altogether, but that is another story.
[4] See, for example, Huemer 1997, Cross 1999, Shogenji 1999, Bovens & Hartmann 2003a, 2003b, 2005, Olsson 2002, 2005, Fitelson 2003, Meijs, 2004, Glass 2006.

which is 'truth-conducive'—that is, whether a higher degree of coherence among a set of propositions ensures that those propositions are more likely to be true, *ceteris paribus*. The general consensus among Bayesian epistemologists is that no probabilistic measure of coherence fully succeeds in being truth-conducive, and this pessimistic consensus is based largely on results by Luc Bovens and Stephan Hartmann (2003a) and Erik Olsson (2005) that show in effect how any probabilistic measure of coherence will fail to ensure a corresponding 'boost' in their likelihood of truth.

The question is whether there are relationships between probabilistically correlated evidence (thought to model 'coherence') and incremental confirmation (thought to model 'justification'), and Bayesian epistemology has investigated this relationship in terms of *models for witness testimony* (Olsson 2002, 2005, Bovens and Hartmann 2003a, 2003b). Think of a witness model as composed of two sets of things, a group of messengers, each with a message to deliver, and the contents of those messages, which we may gather together to form an *information set*. Olsson's model differs from Bovens and Hartmann's model in important ways, but both share two key assumptions about the structure of a *Bayesian witness model*: namely

(bw1) a messenger i who reports that A is true is positive evidence for A, that is, $\Pr(A \mid \text{Report}_i(A)) > \Pr(A)$, and
(bw2) that each messenger is an independent reporter, that is, whether A or ¬A *screens off*[5] whether messenger i reports A or reports ¬A from all other contingent facts and all other messenger reports.

The idea is to ensure that a messenger considering whether to report A or its negation is only influenced by whether in fact A or its negation are true, and not what other facts might be or what other messengers might be saying.

According to Olsson, these two assumptions—the twin pillars of Bayesian witness models—offer not only the most favorable circumstance in which to see if there is any hope of showing that some probabilistic measure of coherence can possibly be truth-conducive, but necessary conditions as well:

> …coherence cannot be truth conducive in the comparative sense in the absence of independence and individual credibility (2005, p. 3).
>
> While these assumptions may seem restrictive from a formal perspective, they should…in fact be seen as describing fortunate circumstances…[and what Olsson's impossibility] theorem says is that not even under fortunate circumstances can there be any interesting measure of coherence or agreement that is truth conducive in the comparative sense (2005, p. 135).

---

[5] See (Pearl 2000) and (Spirtes *et*. *al*. 2000) for a thorough treatment.

And yet, while (bw1) and (bw2) may seem intuitively both favorable and necessary, it turns out neither is the case. The witness testimony models are among the least favorable models for exploring the relationship between coherence and likelihood of truth (Wheeler and Scheines 2011 forthcoming, propositions 3 and 4). The problem is the Bayesian witness model, not the general features of measures of coherence and confirmation.

Indeed, if you drop the conditional independence condition that is built into Bayesian witness models—assumption (bw2) above—there is a measure of association, called *focused correlation* (Myrvold 1995, Wheeler 2009), which robustly tracks incremental confirmation (Wheeler and Scheines 2011, Schlosshauer and Wheeler 2011, Wheeler and Scheines, forthcoming).

Briefly, focused correlation is the ratio of two quantities, the degree of association among evidence given a hypothesis, over the degree of association in the evidence alone. This relationship is clearest in the rightmost expansion of the measure, *For*, which is defined here for two evidence statements, *E*1 and *E*2, and a single hypothesis, *H*.

$$For_H(E_1, E_2) =_{df} \frac{\Pr(H \mid E_1, E_2)}{\Pr(H \mid E_1)\Pr(H \mid E_2)} = \frac{\dfrac{\Pr(E_1, E_2 \mid H)}{\Pr(E_1 \mid H)\Pr(E_2 \mid H)}}{\dfrac{\Pr(E_1, E_2)}{\Pr(E_1)\Pr(E_2)}}.$$

Given some provisos,[6] it turns out that comparing one evidence set (e.g., $\{E_1, E_2\}$) to another (e.g., $\{E_1, E_3\}$) by their degree of focused correlation (with respect to a designated hypothesis *H*), more focused correlation entails more incremental confirmation, *ceteris paribus* (Wheeler and Scheines forthcoming, Schlosshauer and Wheeler, 2011). What's more, making even weaker assumptions, when focused correlation of an evidence set (with respect to a hypothesis, *H*) is greater than 1, then the incremental confirmation of *H* given that evidence is positive (Wheeler 2009, Wheeler and Scheines, forthcoming).

Our point is not that the Bayesian impossibility results fail to be theorems, but rather that they are only representative of a narrow class of models. What we see is that while the tracking condition for focused correlation doesn't hold within witness models—because *ceteris* is not *paribus*—it works fine in many cases outside of the constraints imposed by the witness models. Why? The reason boils down to an insight from focused correlation, which is that there is a parameter missing from previous attempts to give a probabilistic theory of coherence. For it is not enough to look at the association of evidence (e.g., the event of messengers all telling a similar story), but rather we must account for the *reason* for that association. After all, witnesses might agree to agree without any regard to the truth of the matter. This is

[6] Conditions (A1, A2) in (Wheeler and Scheines forthcoming) which is generalized in (Schlosshauer and Wheeler 2011)

the possibility that (bw2) was designed to prevent. But, in so doing, the witness models also inadvertently scupper the possibility of connecting a higher measure of association to higher likelihood of truth. That is the surprising and important insight from the Bayesian impossibility results.

Given this observation that one must account for the cause of the association ('coherence'), Scheines and I have proposed a general model which takes account of the causal structure regulating the relationships between the evidence (or messenger reports, if you prefer) and hypothesis (the truth of the reports, if you prefer). This requires rethinking the commitment to Lewis-BonJour witness models, and moving away from defining coherence solely in terms of information sets, but BonJour himself seems to have already done this:

> 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 justificationally irrelevant), appeals instead to a general belief that beliefs caused in this special way (and perhaps satisfying further conditions as well) are generally true (2002, p. 206-7).

In summary, we think that taking into account what causes the coherence is crucial to making progress on a formal, probabilistic theory of coherence, and we've followed through on this idea by introducing a model which combines focused correlation, for cases in which there are no independence conditions to foul up the probabilistic machinery, and causal structure to help identify how associated evidence will affect incremental confirmation. The model is hardly comprehensive, and there are intriguing irregularities which the impossibility results allude to. (Indeed, we present a more general version of Olsson's impossibility result in Wheeler and Scheines forthcoming).) But, we think that it is a step in the right direction for solving the riddle of coherence.

Dilating sets of probabilities

Open an introductory textbook on probability and within the first few pages you will invariably find a definition of stochastic independence. Defined with respect to a classical probability function, Pr, we say that event E is *stochastically independent* of event F just in case the joint probability distribution of both E and F is equal to the product of the marginal distribution for E and the marginal distribution for F, that is:

(IND) $\Pr(E,F) = \Pr(E) \times \Pr(F)$.

For example, suppose that E is the event of a fairly flipped 1 Euro coin landing 'tails' and F is the event of a fairly flipped American quarter landing 'tails'. The two tosses are stochastically independent just when the probability of both coins landing 'tails' is ¼.

Those textbooks often will give an alternative definition, too. So long as $\Pr(F)$ is non-zero, we may also say that E is stochastically independent of F just when F is *epistemically irrelevant* to the probability of E:

(IR) $\Pr(E|F) = \Pr(E)$, when $\Pr(F) > 0$.

Event F is epistemically irrelevant to E when there is no difference between the probability of E conditional on F and the probability of E alone. Returning to our coins, the probability that a fairly tossed American quarter landing 'tails' given that a fairly tossed 1 Euro coin has landed 'tails' is ½, which is the same as the probability that a fairly tossed American quarter lands 'tails'. In other words, knowing how the experiment with the Euro turns out is irrelevant to estimating the outcome of a fairly tossed quarter.

Finally, we may just as well switch the places of F and E—so long as we make the appropriate accommodations to avoid conditioning on zero-probability events. Let's say then that E is *epistemically independent* of F just when each is epistemically irrelevant to the other, that is:

(EI) $\Pr(E \mid F) = \Pr(E)$, when $\Pr(F) > 0$ and $\Pr(F \mid E) = \Pr(F)$, when $\Pr(E) > 0$.

When working with a single probability distribution, Pr, these notions are equivalent when $\Pr(F) > 0$ and $\Pr(E) > 0$: that is, (IND) iff (IR) iff (EI). Indeed, you are unlikely to see names for each of these notions in your textbook, since they are generally thought to be expressions of one and the same concept: *probabilistic independence*.[7]

However, there are differences between each of these three notions. If Pr is given a behavioral interpretation and viewed to represent an agent's degrees of belief, then arguably that agent learns that two events are stochastically independent (IND) by observing that one event is epistemically irrelevant to the other (IR). In other words, on a behavioral interpretation of Pr, the way for an agent to justify that two events are stochastically independent is from observing that one event is epistemically irrelevant to the other. Furthermore, since we know that (IR) if and only if (EI), the notion of epistemic independence seems an unnecessary, intermediary step. Finally, knowing that a joint distribution satisfies (IND) licenses

[7] Or if there is a distinction draw, it is simply between conditional independence (IR) and independence (IND).

us to factorize that joint distribution by the marginal distributions, which gives probability some semblance of acting like a logic since the probability of (E and F) is determined by taking the product of the probability of E and the probability of F. The ability to factorize a joint distribution into a product of marginal probabilities and conditional probabilities is a tremendous advantage to computing probabilities.

So, one way to look at the equivalence of (IND), (IR), and (IND), is that it licenses learning about stochastic independence through observing when one event is irrelevant to the probability estimate of another, and then allows us to leverage what we learn about those independence conditions to yield tractable methods for probabilistic reasoning. The theory of causal Bayes nets illustrates this strategy perfectly (Pearl 2000, Spirtes *et. al*. 2000). Our own approach to probabilistic logic is another example (Haenni et. al. 2011).

Yet, there is a question of how sound a foundation this strategy rests on. It turns out, surprisingly, that these three independence concepts are distinct mathematical notions after all. They only appear to be three equivalent ways of expressing the same concept when viewed through the lens of a single probability measure, Pr. What this means, philosophically, is that sound probabilistic reasoning from independence and conditional independence assumptions which glide freely between (IND), (IR), and (EI), will turn out to depend on reasoning with a single probability distribution. So, those sounds principles of probabilistic reasoning depend on the grounds that you have for assuming a numerically precise probability distribution.

Assuming that agents *always* have numerically determinate degrees of belief, however, is a stretch, and several authors have argued that probability models should accommodate approximate or interval values.[8] Take for example that coin in your pocket: it is an idealization to assume that the probability of a single, fair toss of that coin landing 'tails' is precisely ½. Instead, the argument for imprecision goes, it is more reasonable to assume that the chance of 'tails' is ½ plus or minus some small ε. Yet opening the door to interval-valued probabilities even just a crack introductions a number of difficult issues, some of which go to the heart of probabilistic reasoning. The question of whether there is one concept of probabilistic independence or several independence concepts is an example.

Suppose that the probability of E is the interval $[l,u]$, where l is understood to be the lower bound of the probability that E, and u is the upper bound of the probability that E. There are several distinct ways to flesh this idea out,[9] but a

[8] Pioneers of imprecise probability theory includes B.O. Koopman (1940), Alfred Horn and Alfred Tarski (1948), Paul Halmos (1950), I. J. Good (1952), C.A.B. Smith (1961), Daniel Ellsberg (1961), and Henry Kyburg, Jr. (1961). Notable contemporary advocates include Isaac Levi (1980), Peter Walley (1991), Teddy Seidenfeld (2010), James Joyce (2010), Fabio Cozman (2000), Gert de Cooman and Enrique Miranda (2007, 2009). See also Haenni et. al. (2011).
[9] Recent textbook treatments include Paris 1994, Halpern 2003, Haenni et. al. 2011.

common one is to interpret an interval probability assignment $[l,u]$ to an event E by a set of probability functions, $\mathbf{Pr} = \{\mathrm{Pr}_1, \mathrm{Pr}_2,\ldots, \mathrm{Pr}_n\}$, where the *lower probability* of E is the *infimum* of $\mathbf{Pr}(E)$, and the *upper probability* of E is the *supremum* of $\mathbf{Pr}(E)$. Define lower and upper probability, with respect to a set of probabilities **Pr**, as:

$$\underline{\mathrm{Pr}}(E) = l = \inf_{\mathrm{Pr}\in\mathbf{Pr}} \mathrm{Pr}(E) \quad \text{(lower probability)}$$
$$\overline{\mathrm{Pr}}(E) = u = \sup_{\mathrm{Pr}\in\mathbf{Pr}} \mathrm{Pr}(E) \quad \text{(upper probability)}$$

As it stands, the set **Pr** could represent a set of classical Bayes agents who express different judgments about E; **Pr** could also be interpreted as a model for studying sensitivity and robustness in classical Bayesian statistical inference; and **Pr** can also be viewed as a model of imprecise credal probabilities for a single agent. This is the interpretation which has drawn fire from epistemologists, but it is worth pointing out that the underlying mechanics of the example we're about to consider is not tied to this interpretation.

Although imprecise probability theory is sometimes described as 'exotic', in reality classical Bayesianism drops out as a special case: when the set **Pr** contains just one measure. From what we've said, the set **Pr** may contain a single measure or two or several different measures. If classical Bayesianism marks one end of set-based Bayesianism, then *convex Bayesianism* (Levi 1980) marks another important endpoint. For we can think of convexity as a closure condition on the set **Pr**:

(Cx) **Pr** is *a closed convex set* when, for any two probability measures Pr_1, Pr_2 in **Pr**, then for all $0 \leq r \leq 1$, the measure $\mathrm{Pr}^* = r\,\mathrm{Pr}_1 + (1-r)\mathrm{Pr}_2$ is also in **Pr**.

Condition (Cx) says that, for any two measures in **Pr**, the measure Pr* defined by the convex mixture of those two measures is also in **Pr**. Informally, we say that a set satisfying (Cx) is closed under convex mixtures, and adding (Cx) as a condition is common when a set of probabilities is interpreted to represent imprecise credences. In what follows, we will assume (Cx) holds. With these preliminaries in place, let's turn to dilation.

We say that an event F *dilates* the event E just in case

$$\underline{\mathrm{Pr}}(E \mid F) < \underline{\mathrm{Pr}}(E) \leq \overline{\mathrm{Pr}}(E) < \overline{\mathrm{Pr}}(E \mid F).$$

In words, outcome F dilates E just in case the range of unconditional probability assignments to E is a proper subset of the range of probability assignments to E given F. Now suppose that **B** is a measurable partition of possible outcomes. Then, the partition of outcomes **B** *strictly dilates* E just in case:

$$\underline{\mathrm{Pr}}(E \mid F) < \underline{\mathrm{Pr}}(E) \leq \overline{\mathrm{Pr}}(E) < \overline{\mathrm{Pr}}(E \mid F), \quad \text{for all } F \in \mathbf{B}.$$

The remarkable thing about strict dilation is the specter of turning a more precise estimate of E into a less precise estimate, *no matter the outcome*.

To illustrate strict dilation, we recount Peter Walley's canonical coin tossing example.

> **Example**. Suppose that a fair coin is tossed twice. The first toss of the coin is a fair toss, but the second toss is performed in such a way that the outcome may depend on the outcome of the first toss. Nothing is known about the type or degree of the possible dependence. Let H_1, T_1, H_2, T_2 denote the possible outcomes for the pair of tosses. We know the coin is fair and that the first toss is a fair toss, so the Agent's ($\Pr_A$) estimate for the first toss is precise. The interaction between the tosses is unknown, but in the extreme the first toss may determine the outcome of the second. This likewise puts a precise constraint on A's estimate of the second toss prior to the experiment. Hence,
>
> $$(a)\quad \underline{\Pr}_A(H_1) = \overline{\Pr}_A(H_1) = \Pr_A(H_1) = \frac{1}{2} = \underline{\Pr}_A(H_2) = \overline{\Pr}_A(H_2) = \Pr_A(H_2).$$
>
> However, little is known about the direction or degree of dependence between the pair of tosses. Model A's ignorance by
>
> $$(b)\quad \underline{\Pr}_A(H_1,H_2) = 0, \text{ and } \overline{\Pr}_A(H_1,H_2) = \Pr_A(H_1) = \frac{1}{2}$$
>
> Suppose now that A learns that the outcome of the first toss is *heads*. The extremal points from (b), namely 0 and ½, can be conditioned by Bayes' rule yielding
>
> $$(c)\quad (i)\ \underline{\Pr}_A(H_2 \mid H_1) = \underline{\Pr}_A(H_2,H_1)/\underline{\Pr}_A(H_1) = 0 \text{ and}$$
> $$(ii)\ \overline{\Pr}_A(H_2 \mid H_1) = \overline{\Pr}_A(H_2,H_1)/\overline{\Pr}_A(H_1) = 1$$
>
> So, although initially $\Pr_A(H_2)$ = ½, learning that the first toss lands heads dilates A's estimate of the second toss to any value within the interval [0,1]. An analogous argument holds if instead A learns that the outcome of the first toss is tails. Since these two outcomes partition the outcome space, i.e., there are no other ways the first toss can turn out, A's precise probability about the second toss strictly dilates to the vacuous unit interval, no matter which way the first coin toss lands (Walley 1991, pp. 298-9).

One way to interpret the two extreme points is that they stand for two opposing hypotheses about the mechanism controlling the second toss. Each hypothesis specifies a deterministic mechanism: case (*i*) says that the second coin is certain to be tails if the first is heads, whereas case (*ii*) says that the second coin is certain to be heads if the first coin toss is heads. So, on this interpretation, the agent knows that the outcome of the first toss may provide relevant information, possibly definitive information, but merely observing that outcome is insufficient to determine in what way the information is relevant. Arguably, then, the outcome of the first toss gives the agent information—or, better, signals a potentially important gap in her information—which warrants the change in belief. This is by no means universally accepted, but it is a more plausible position than the next example, which seems to reproduce the same result while assuming that there is no connection between the first toss and the second toss.

Suppose that instead of tossing a single coin twice and hiding from the agent the manner in which the second toss is performed, we instead flip two different coins, one which is known to be normal and the other of unknown bias. Whereas in the first coin example the methods used for performing the tosses varied but the constitution of the coin remained fixed, in this example the mechanism for tossing the coins is normal—that is, the coin tosses are independent—but the consitution of the coins are different.[10] What we are doing is replacing (a) in the example above by

$$(a') \quad \underline{\Pr}_A(H_2) = \overline{\Pr}_A(H_2) = \Pr_A(H_2) = \frac{1}{2}$$
$$\underline{\Pr}_A(H_1) = 0, \overline{\Pr}_A(H_1) = 1,$$

which appears to dilate the second toss as well. But the explanation we provided for the first example is not available to explain this second example, since we have stipulated that the tosses are independent. Yet, even though the two events are independent, it appears that the imprecision of one event can dilate the sharp probability estimate of another, independent event. How can this be?!

What is interesting is that 'independent event' in this setting is ambiguous between analogues of *stochastic independence*, *epistemic independence*, and *epistemic irrelevance* which are defined for sets of probabilities. A necessary condition for dilation is for E and F to not be stochastically independent (Seidenfeld et. al. 1992, Theorems 2.1-2.3), which is a bulwark against the paradoxical conclusions critics have drawn. What Seindenfeld et. al.'s theorems tell us is that where there is dilation, there is dependence. So, it would appear, the flips in the second coin toss are dependent after all. How could this be?

The answer returns us to our earlier discussion of independence: there are several independence concepts rather than a single, unified independence concept (Kyburg and Pittarelli, Cozman 2012, Wheeler, forthcoming). Within the imprecise probability setting, and assuming (Cx), stochastic independence entails epistemic independence, and epistemic independence entails epistemic irrelevance, but it is not the case that epistemic irrelevance entails epistemic independence nor, shockingly, does epistemic independence entail stochastic independence! The

---

[10] A version of this is discussed by Walley (1991), White (2010), Sturgeon (2010), and Joyce (2010).

reasons for this are technical, and will have to be dealt with in another essay.[11] But the larger, philosophical point is that recent discussions about dilation have foundered on a real (if understandable) fallacy in probabilistic reasoning which hinge on assuming, falsely, that probabilistic independence is a unitary notion, and that the principles of sound probabilistic reasoning which hold for a single measure extend to settings in which there are sets of measures.

This observation suggests the controversy over dilation is a side issue, and that really the issue is that there are a plurality of independence concepts. This points to a dilemma for Bayesianism, with orthodox Bayesians on one horn, and set-based Bayesians on the other:

- ***For orthodox Bayesians***: Imprecise probability theory reveals a fact about (IND), (EI), and (IR), namely that they are distinct properties which are collapsed when working with a single probability distribution, Pr. Orthodox Bayesianism hides this fact from view. However, even if you reject set-based approaches, are you confident that your elicitation procedure for determining numerically precise degrees of belief warrants collapsing these distinctions?

- ***For set-based Bayesians***: In so far as you rely on a behavior interpretation of your convex set of distributions, how do you provide a *behavioral* justification for treating two events as completely stochastically independent given that (IND) does not follow from (EI)?

In short, the discovery that there are many independence concepts rather than a single concept is an example of research within formal epistemology that has far reaching consequences. Although we have illustrated this discovery through diagnosing a common missteps in the recent literature on dilation, the underlying point concerns the very foundations of sound probabilistic reasoning.

The FIE-model of inquiry

C.P. Snow observed long ago that universities are made up of two broad types of people, literary intellectuals and hard scientists, yet a typical individual of each type is barely able, if able at all, to communicate with his counterpart. Snow's observation, popularized in his 1959 lecture *Two Cultures and the Scientific Revolution* (reissued by Cambridge 1993), goes some way to explaining the two distinct cultures one hears referred to as 'the humanities' and 'the sciences'.

Certainly there is some basis for grouping academic subjects the way we do. Physics,

---

[11] However, if you drop (Cx), then Stochastic independence does not entail Epistemic independence. See Wheeler forthcoming for an example.

chemistry, and biology are the pillars of experimental science. Although the skills and methods differ from each, all aim to reconcile theory purporting to govern some part of the natural world against the evidence. However, the subjects studied by the humanities typically don't yield to experimental data; there are no experimental branches of history, no laboratories of literature. It is tempting to view the importance placed on experimental data as an indication of how scientific a subject is. The hard sciences put experimental evidence front and center, whereas the humanities either do not or cannot. The quarrels over experimental philosophy, and to a some extent, debate over formal epistemology, are very often viewed as pitched battles about whether philosophy is part of the humanities or part of the sciences. Heaped on top of that quarrel is another about whether philosophy even should be placed in one or the other.

Although this is a familiar picture, it is misleading. Mathematics has no more to do with experimental data than poetry, and professional cooking is as concerned with experimentation as any of the social sciences. But cooking is clearly a trade, not an academic subject, much less a science.

In closing, I want to suggest that we should instead think of academic disciplines as dividing into three categories rather than two. There are *formal* disciplines, *experimental* disciplines, and *interpretive* disciplines. This three-way distinction was proposed by Henry Kyburg in *Science and Reason* (1990, 16) to better represent the activities that make up a university, but there is much to recommend this way of thinking about academic disciplines, particular those—like philosophy—which are restlessly interdisciplinary in nature. Call this the FIE-model of inquiry.

Mathematics is essentially a formal discipline, the empirical sciences are largely empirical disciplines, and the traditional fine arts and letters are the leading exemplars of the interpretive disciplines. But nearly all fields draw upon skills from each category. Biology and literature are often concerned with formal structures, mathematics and psychology are sometimes concerned with interpretation, and psychology and literature are at various times interested in the facts about the society or groups that produced an observed behavior, or whose members wrote a series of plays.

It is unclear whether this FIE-model would help in organizing a university, as Kyburg suggested when he proposed it. The idea would need an administrative Helmholtz to sort that question out. Or, perhaps, a Clausewitz. But, the categories are helpful for a scholar to have in mind when working on a topic like *epistemology*. It is not so much that epistemology calls upon results from the cognitive, computational, and decision sciences, and insights from philosophy, mathematics, psychology, statistics, and linguistics—although it certainly does. Rather, the central point is that epistemology calls upon the full range of inquiry, even if truly harnessing together formal, experimental, and interpretive *skills* do not readily match the way we have happened to organize our departments and our universities. That Snow was right is surely no reason to believe that he must go on being right.

References

BonJour, Laurence 1985: *The Structure of Empirical Knowledge*. Cambridge, MA, Harvard University Press.

BonJour, Laurence 1999: 'The Dialectics of Foundationalism and Coherentism'. In Greco and Sosa 1999, pp. 117-42.

BonJour, Laurence 2002: *Epistemology*. Oxford: Rowman and Littlefield.

Bovens, Luc and Stephan Hartmann 2003a: *Bayesian Epistemology*. Oxford: Oxford University Press

Bovens, Luc and Stephan Hartmann 2003b: 'Solving the Riddle of Coherence'. *Mind*, 112, pp. 601-33.

Bovens, Luc and Stephan Hartmann 2006: 'An Impossibility Result for Coherence Rankings'. *Philosophical Studies*, 128, pp. 77-91.

Bovens, Luc and Erik Olsson 2000: 'Coherentism, Reliability, and Bayesian Networks'. *Mind*, 109: 685-719.

Breese, J. and D. Koller (eds) 2001: *Uncertainty in artificial intelligence: Proceedings of the 17th conference (UAI-2001)*. San Francisco: Morgan Kaufmann.

Carlson, Rae: 1971: 'Where is the person in personality research?' *Psychological Bulletin*, 75, p. 212.

Carnap, Rudolf 1962: *The Logical Foundations of Probability*. Chicago: University of Chicago Press.

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

Cozman, Fabio 2000: 'Credal Networks', *Artificial Intelligence*, 120(2), pp. 199-233.

Cozman, Favio 2012: 'Sets of Probability Distributions, Independence, and Convexity'. *Synthese*, in press.

Cross, Charles B. 1999: 'Coherence and Truth Conducive Justification'. *Analysis*, 59(3), pp. 186-93.

Crupi, V., K. Tentori, and M. Gonzalez 2007: 'On Bayesian Measures of Evidential Support: Theoretical and empirical issues'. *Philosophy of Science*, 74(2), pp. 229-52.

Danks, David and Clark Glymour 2001: 'Linearity Properties of Bayes Nets with Binary Variables'. In Breese and Koller 2001, pp. 98-104.

de Cooman, Gert and Enrique Miranda 2007: 'Symmetry of models versus models of symmetry'. In Harper and Wheeler 2007, pp. 67-149.

de Cooman, Gert and Enrique Miranda 2009: 'Forward Irrelevance'. *Journal of Statistical Planning*, 139, pp. 256-76.

Dennis, Wayne (ed) 1948: *Readings in the History of Psychology*. East Norwalk, CT: Appleton, Century, Crofts.

Douven, Igor and Wouter Meijs 2007: 'Measuring Coherence'. *Synthese,* 156(3), pp. 405-25.

Earman, John 1992: *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.

Eells, Ellery and Branden Fitelson 2002: 'Symmetries and Asymmetries in Evidential Support'. *Philosophical Studies*, 107(2), pp. 129-42.

Ellsberg, Daniel 1961: 'Risk, Ambiguity, and the Savage Axioms'. Quarterly Journal of Economics, 75: pp. 643-69.

Evans, J. St. B. T. 2008: 'Dual-processing Accounts of Reasoning, Judgement, and Social Cognition'. *Annual Review of Psychology*, 59, pp. 255-78.

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

Fitelson, Branden 2003: 'A Probabilistic Theory of Coherence'. *Analysis*, 63, pp. 194-99.

Gendler, Tamar and John Hawthorne (eds) 2010: *Oxford Studies in Epistemology, volume 3*. Oxford: Oxford University Press.

Gettier, Edmund 1963: 'Is Justified True Belief Knowledge?' *Analysis*, 23, pp. 121-123.

Good, I. J. 1952: 'Rational Decisions'. *Journal of the Royal Statistics Society, Series B*. 14(1), pp. 107-14.

Greco, John and Ernest Sosa (eds) 1999: *The Blackwell Guide to Epistemology*. Malden, MA: Blackwell.

Glass, D. H. 2006: 'Coherence Measures and their Relations to Fuzzy Similarity and Inconsistency in Knowledge Bases'. *Artificial Intelligence Review*. 26, pp. 227-49.

Glymour, Clark 1998: 'What Went Wrong: Reflections on Science by Observation and The Bell Curve'. *Philosophy of Science*, 65(1), pp. 1-32.

Glymour, C., R. Scheines, P. Spirtes, and K. Kelly 1987: *Discovering Causal Structure*. London: Academic Press.

Haenni, R. J.W. Romeyn, G. Wheeler, and J. Williamson 2011: *Probabilistic Logic and Probabilistic Networks*, Dordrecht: The Synthese Library.

Halmos, Paul 1950: *Measure Theory*. New York: Van Nostrand Reinhold Company.

Halpern, Joseph 2003: *Reasoning about Uncertainty*. Cambridge, MA: MIT Press.

Harper, William and Gregory Wheeler (eds) 2007: *Probability and Inference: Essays in Honour of Henry E. Kyburg, Jr.* London: College Publications.

Hartmann, Stephan, Marcel Weber, Wenceslao J. Gonzalez, Dennis Dieks, Thomas Uebe (eds) 2011: *Explanation, Prediction, and Confirmation: New Trends and Old Ones Reconsidered*. Dordrecht: Springer.

Helmholtz, Hermann 1860: 'The Young-Helmholtz Theory of Color Vision', in Dennis, Wayne 1948, pp.199-205.

Hendricks, Vincent and John Symons 2005: *Formal Philosophy*. Rolskilde: Automatic Press.

Henrich, Joseph, Steven Heine and Ara Norenzayan 2010: The Weirdest People in the World? *Behavioral and Brain Sciences*. 33(2-3), pp. 61-83.

Horn, Alfred and Alfred Tarsk 1948: 'Measures in Boolean Algebras'. *Transactions of the AMS*. 64(1): 467-97.

Huemer, Michael 1997: 'Probability and Coherence Justification'. *The Southern Journal of Philosophy*, 35, pp. 463-72.

Jeffrey, Richard 1965: *The Logic of Decision*. New York: McGraw-Hill.

Joyce, James 2010: 'In Defense of Imprecise Credences in Inference and Decision Making'. *Philosophical Perspectives*, 21(1): 281-323.

Junker, B. W. and J. L. Ellis 1997: 'A Characterization of Monotone Unidimensional Latent Variable Models'. *The Annals of Statistics*, 25, pp. 1327-43.

Klein, Peter and Ted Warfield 1994: 'What Price Coherence?' *Analysis*, 54(3), pp.129-32.

Koopman, B. O. 1940: 'The axioms and algebra of intuitive probability'. *Annals of Mathematics*, 41(2): 269-92.

Kruglanski, Arie and Gerd Gigerenzer 2011: 'Intuitive and Deliberate Judgments Are Based on Common Principles'. *Psychological Review*, 118(1): 97-109.

Kruglanski, Arie and Edward Orehek 2007: 'Partitioning the Domain of Human Inference: Dual mode and system models and their alternatives'. *Annual Review of Psychology*, 8, 291-316.

Kyburg, Jr. H. E. 1961*: Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.

Kyburg, Jr. H. E. 1990: *Science and Reason*. Oxford: Oxford University Press.

Kyburg, Jr., H. E. and M. Pittarelli 1996: 'Set-based Bayesianism'. *IEEE Transactions on Systems, Man, and Cybernetics A*. 26(3): 324-39.

Levi, Isaac 1980: *Enterprise of Knowledge*. Cambridge, MA: MIT Press.

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

Meijs, Wouter 2004: 'A Corrective to Bovens and Hartmann's Measure of Coherence'. *Philosophical Studies*. 133(2), pp. 151-80.

Myrvold, Wayne 1996: 'Bayesianism and Diverse Evidence: A Reply to Andrew Wayne'. *Philosophy of Science*. 63, pp. 661-5.

Olsson, Erik J. 2002: 'What is the Problem of Coherence and Truth?' *Journal of Philosophy*. 94, pp. 246-72.

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

Paris, Jeff 1994: *The Uncertain Reasoner's Companion: A Mathematical Companion*. Cambridge: Cambridge University Press.

Sears, D. O. 1986: 'College Sophomores in the laboratory: influences of a narrow data base on social psychology's view of human nature'. *Journal of Personality and Social Psychology*, 51, pp. 515-30.

Schlosshauer, Maximillian and Gregory Wheeler 2011: 'Focused Correlation, Confirmation, and the Jigsaw Puzzle of Variable Evidence'. *Philosophy of Science*, 78(3), pp. 276-92.

Segerberg, Krister 2005: 'Krister Segerberg', in Hendricks and Symons 2005, pp. 159-167.

Seidenfeld, Teddy, Mark J. Schervish and Joseph B. Kadane 2010: 'Coherent Choice Functions under Uncertainty'. Synthese, 172(1), pp. 157-76.

Seidenfeld, Teddy and Larry Wassermann 1993: 'Dilation for sets of probabilities'. *The Annals of Statistics*, 21: 1139-54.

Shogenji, Tomoji 1999: 'Is Coherence Truth Conducive?' *Analysis*, 59, 1999, 338-45.

Silva, R., C. Glymour, R. Scheines, and P. Spirtes 2006: 'Learning the Structure of Latent Linear Structure Models'. *Journal of Machine Learning Research*, 7, pp. 191-246.

Snow, C. P. 1959 (1998): *The Two Cultures*. 9th *Printing.* Cambridge: Cambridge University Press.

Spirtes, P., C. Glymour, and R. Scheines 2000: *Causation, Prediction, and Search*. 2nd edition. Cambridge, MA: MIT Press.

Sturgeon, S. 2010: 'Confidence and Coarse-grain Attitudes'. In Gendler and Hawthorne 2010, pp. 126-49.

Walley, Peter 1991: *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.

Weinberg, J., S. Nichols and S. Stich 2001: 'Normativity and Epistemic Intutions'. *Philosophical Topics*, 29, pp. 429-60.

Wheeler, Gregory 2009: 'Focused Correlation and Confirmation'*. The British Journal for the Philosophy of Science*. 60(1), pp. 79-100.

Wheeler, Gregory: 'Objective Bayesian Calibration and the Problem of Non-convex Evidence'. *The British Journal for the Philosophy of Science*, to appear.

Wheeler, Gregory and Luís Moniz Pereira 2008: 'Methodological Naturalism and Epistemic Internalism'. Synthese, 163(3): 315-28.

Wheeler, Gregory and Richard Scheines 2011: 'Coherence, Association, and Causation'. In Hartmann, Weber, Gonzalez, Dieks, and Uebe, pp. 37-51.

Wheeler, Gregory and Richard Scheines: 'Coherence and Confirmation through Causation'. under review.

White, Roger 2010: 'Evidential symmetry and mushy credence'. In Gendler and Hawthorne 2010, pp. 161-81.

Williamson, Jon 2010: *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.