

A Bayesian Account of Independent Evidence with Applications

Branden Fitelson^{†‡}

University of Wisconsin–Madison

August 15, 2000

A Bayesian account of independent evidential support is outlined. This account is partly inspired by the work of C.S. Peirce. I show that a large class of quantitative Bayesian measures of confirmation satisfy the basic desiderata laid down by Peirce for adequate accounts of independent evidence. I argue that, by considering further natural constraints on a probabilistic account of independent evidence, all but a very small class of Bayesian measures of confirmation can be ruled out. In closing, another application of my account to the problem of evidential diversity is also discussed.

[†]Philosophy Department, University of Wisconsin, 600 N. Park Street, Madison, WI 53706.

[‡]Thanks to Marty Barrett, Ellery Eells, Malcolm Forster, Mike Kruse, Greg Mougin, Dan Steel, and, especially, Patrick Maher and Elliott Sober for useful suggestions.

1 Terminology, Notation & Basic Assumptions

The present paper is concerned with the degree of incremental confirmation provided by evidential propositions E for hypotheses under test H , given background evidence K , according to relevance measures of degree of confirmation \mathfrak{c} . We say that \mathfrak{c} is a *relevance measure* of degree of confirmation if and only if \mathfrak{c} satisfies the following constraints, in cases where E confirms, disconfirms, or is confirmationally irrelevant to H , given background evidence K .

$$(\mathcal{R}) \quad \mathfrak{c}(H, E | K) \begin{cases} > 0 & \text{if } \Pr(H | E \& K) > \Pr(H | K), \\ < 0 & \text{if } \Pr(H | E \& K) < \Pr(H | K), \\ = 0 & \text{if } \Pr(H | E \& K) = \Pr(H | K). \end{cases}$$

I will restrict my attention to the following four relevance measures of degree of confirmation: the *difference* measure d , the *log-ratio* measure r , the *log-likelihood ratio* measure l , and Christensen's (1999) "normalized" difference measure s . These measures are representative of the varieties of quantitative Bayesian confirmation theory that are currently defended in the philosophical literature.¹ The measures d , r , l , and s are defined as follows:²

$$d(H, E | K) =_{df} \Pr(H | E \& K) - \Pr(H | K)$$

$$r(H, E | K) =_{df} \log \left[\frac{\Pr(H | E \& K)}{\Pr(H | K)} \right]$$

$$l(H, E | K) =_{df} \log \left[\frac{\Pr(E | H \& K)}{\Pr(E | \bar{H} \& K)} \right]$$

$$\begin{aligned} s(H, E | K) &=_{df} \Pr(H | E \& K) - \Pr(H | \bar{E} \& K) \\ &= \frac{1}{\Pr(\bar{E} | K)} \cdot d(H, E | K).^3 \end{aligned}$$

¹Many relevance measures have been proposed over the years. For a nice survey, see Kyburg 1983. The three relevance measures d , r , and l have had the most loyal following in recent years. Advocates of d include Earman (1992), Eells (1982), Gillies (1986), Jeffrey (1992), and Rosenkrantz (1994). Advocates of r (or measures ordinally equivalent to r) include Horwich (1982), Keynes (1921), Mackie (1969), Milne (1996), Schlesinger (1995), and Pollard (1999). Advocates of l (or measures ordinally equivalent to l) include Kemeny and Oppenheim (1952), Good (1984), Heckerman (1988), Horvitz and Heckerman (1986), Pearl (1988), and Schum (1994). Recent proponents of s include Christensen (1999) as well as Joyce (1999).

²Overbars are used to express negations of propositions (i.e. ' \bar{X} ' stands for '*not-X*'). Logarithms (of any arbitrary base greater than 1) of the ratios $\Pr(H | E \& K) / \Pr(H | K)$ and $\Pr(E | H \& K) / \Pr(E | \bar{H} \& K)$ are taken to ensure that (i) r and l satisfy \mathcal{R} , and (ii) r and l are *additive* in various ways (see footnote 6). Not all advocates of r or l adopt this convention (e.g., Horwich 1982 and Sober 1989). But, because logarithms are monotonic functions, defining r and l in this way will not result in any loss of (or gain in) generality in my argumentation.

³This equality holds provided, of course, that $\Pr(\bar{E} | K) \neq 0$. See Christensen (1999) for further discussion about the relationship between d and s .

When we want to consider how confirmation varies with changing background evidence, we will use the conditional notation $\mathfrak{c}(H, E_1 | E_2)$ to denote the degree to which E_1 confirms H (according to \mathfrak{c}), given that E_2 is part of our background evidence.⁴ And, we will use the unconditional notation $\mathfrak{c}(H, E_1)$ to denote the degree to which E_1 confirms H (according to \mathfrak{c}), *not* conditional on E_2 being part of our background evidence.

2 Confirmation Independence — Bayesian Style

2.1 The Fundamental Peircean Desiderata

In his essay “The Probability of Induction”, C.S. Peirce articulates several fundamental intuitions concerning the nature of independent inductive support. Consider the following important excerpt from Peirce (1878, my brackets):

... two arguments which are entirely independent, neither weakening nor strengthening the other, ought, when they concur, to produce a[n intensity of] belief equal to the sum of the intensities of belief which either would produce separately.

Two crucial intuitions about independent inductive support are contained in this quote. First, there is the intuition that two pieces of evidence E_1 and E_2 provide *independent* inductive support for a hypothesis H if the degree to which E_1 supports H does not depend on whether E_2 is part of our background evidence (and vice versa). In our confirmation-theoretic framework, we will take this intuition onboard as a *definition* of (mutual) confirmation independence regarding a hypothesis:⁵

Definition. E_1 and E_2 are (mutually) *confirmationally independent* regarding H according to \mathfrak{c} iff both $\mathfrak{c}(H, E_1 | E_2) = \mathfrak{c}(H, E_1)$, and $\mathfrak{c}(H, E_2 | E_1) = \mathfrak{c}(H, E_2)$.

The second intuition expressed by Peirce in this passage is that the joint support provided by two pieces of independent evidence should be *additive*. In our confirmation theoretic framework, this gets unpacked as follows:

(A') If E_1 and E_2 are confirmationally independent regarding H according to \mathfrak{c} , then $\mathfrak{c}(H, E_1 \& E_2) = \mathfrak{c}(H, E_1) + \mathfrak{c}(H, E_2)$.

⁴There may be other background evidence besides E_2 in a confirmational context. However, this additional background evidence will be *held fixed* in the confirmational comparisons we do to determine whether E_1 and E_2 are dependent or independent regarding H in that context. So, there is no need to indicate this additional background evidence explicitly. As such, I will, for simplicity, hereafter suppress the (full) background evidence K from my notation.

⁵Our formal characterization of confirmation independence regarding a hypothesis is very similar to the formal notion of ‘modularity’ in the recent artificial intelligence literature. For a nice survey of some recent applications (and *mis*applications) of the modularity concept in artificial intelligence, see Horvitz and Heckerman (1986), and Heckerman (1988).

Strictly speaking, we should weaken \mathcal{A}' to require only that $\mathfrak{c}(H, E_1 \& E_2)$ be *some* (symmetric) monotonic function f of $\mathfrak{c}(H, E_1)$ and $\mathfrak{c}(H, E_2)$, where f is additive in *some* (monotonically) transformed space.⁶ The point is that, if E_1 and E_2 are confirmationally independent regarding H according to \mathfrak{c} , then $\mathfrak{c}(H, E_1 \& E_2)$ should depend *only* (and, in some appropriately transformed space, *linearly*) on $\mathfrak{c}(H, E_1)$ and $\mathfrak{c}(H, E_2)$, without any extra “interaction terms”. This leads to the following refinement of the second basic Peircean intuition:

- (\mathcal{A}) There exists some (symmetric) monotonic function f such that, for all E_1 , E_2 , and H , if E_1 and E_2 are confirmationally independent regarding H according to \mathfrak{c} , then $\mathfrak{c}(H, E_1 \& E_2) = f[\mathfrak{c}(H, E_1), \mathfrak{c}(H, E_2)]$, where f is additive (*i.e.*, linear) in some (monotonically) transformed space.

The following theorem states that each of our four Bayesian relevance measures — *except s* — satisfies \mathcal{A} (see the Appendix for proofs of all theorems).

Theorem 1. *Each of the measures d , r , and l satisfies \mathcal{A} , but s violates \mathcal{A} .*⁷

So, at this most basic level, the three most popular varieties of quantitative Bayesian confirmation theory are in agreement about the nature of independent evidence. All three measures d , r , and l satisfy the fundamental Peircean desideratum \mathcal{A} (and \mathcal{A}'). However, measure s would seem to be inadequate in its handling of independent evidence, even at this most basic level.⁸

The agreement between d , r , and l ends here. In the next section, I will describe a symmetry desideratum which is satisfied by d and l (and s), but violated by r . This will narrow down the field further to two measures (d and l) which seem to cope adequately (at a very basic level) with independent evidence. Later, I will propose additional, probabilistic constraints on accounts of independent evidence that narrow the field even more.

2.2 A Negation Symmetry Desideratum

If two pieces of evidence are confirmationally independent regarding H , then they should also be confirmationally independent regarding \bar{H} . Negation symmetry in the independence relation seems highly intuitive.⁹ After all, if the degree to which E_1 confirms H doesn't depend on whether E_2 is already known, then why should the degree to which E_1 confirms \bar{H} depend on whether E_2 is already known? In our confirmation theoretic framework, this intuitive negation symmetry principle gets formalized as follows:

⁶As Peirce did, I prefer to have f be $+$. So, I have defined r and l using *logarithms* (see footnote 2). If we were to drop the logarithms in our definitions of r and l , then we would have $f = \cdot$ for the ratio measures r and l , but $f = +$ for the difference measure d . See Heckerman (1988) for more on the kind of linear decomposability that is at the heart of desideratum \mathcal{A} .

⁷Carnap's (1962, §67) relevance measure \mathfrak{r} also violates \mathcal{A} (proof omitted). See Fitelson (1999) and Eells and Fitelson (2000b) for other unintuitive features of Carnap's measure \mathfrak{r} .

⁸See Eells and Fitelson (2000a, 2000b) for other unintuitive features of s .

⁹Many varieties of independence satisfy this kind of negation symmetry requirement (*e.g.*, both logical independence and probabilistic independence are negation-symmetric).

(S) If $c(H, E_1 | E_2) = c(H, E_1)$ and $c(H, E_2 | E_1) = c(H, E_2)$, then
 $c(\bar{H}, E_1 | E_2) = c(\bar{H}, E_1)$ and $c(\bar{H}, E_2 | E_1) = c(\bar{H}, E_2)$.

The following theorem states that each of our four Bayesian relevance measures — *except* r — satisfies \mathcal{S} .

Theorem 2. *Each of the measures d , l , and s satisfies \mathcal{S} , but r violates \mathcal{S} .*¹⁰

The two high-level desiderata \mathcal{A} and \mathcal{S} narrow the field of four relevance measures down to two (d and l) which seem — so far — to explicate the concept of independent evidence.¹¹ Next, I will propose a low-level, probabilistic constraint that rules out the difference measure d and all other relevance measures, except those ordinally equivalent to the log-likelihood ratio measure l .

2.3 Screening-Off and Confirmational Independence

2.3.1 Sober’s Intuitive Conjunctive Fork Example

Sober (1989) discusses an example in which two pieces of evidence E_1 and E_2 seem — intuitively — to provide independent support in favor of a hypothesis.¹² In Sober’s example, E_1 is a newspaper report of the outcome H of a baseball game, and E_2 is an independently derived radio report of the (same) outcome of the same baseball game. It is assumed that each of E_1 and E_2 *individually* confirms H . Sober explains that the (intuitive) probabilistic structure of this example is a *conjunctive fork*, in which E_1 and E_2 are joint effects of a common cause H . Sober also points out (as Reichenbach 1956, page 159 first did) that E_1 and E_2 will *not* be *unconditionally* probabilistically independent in such a case. So, it *can’t* be probabilistic independence of the evidence *simpliciter* which is responsible for our intuitive judgment that E_1 and E_2 are confirmationally independent *regarding* H in Sober’s example. Is there *some* probabilistic feature of Sober’s example which undergirds our intuition? It seems to me (as it did to Sober) that the relevant point is that (in the terminology of Reichenbach 1956, page 189) each of H and \bar{H} *screens-off* E_1 from E_2 . That is, it is the fact that E_1 and E_2 are probabilistically independent *conditional on the hypothesis* H (and its denial) that undergirds our intuition that E_1 and E_2 are *confirmationally independent regarding* H .

Sober’s conjunctive fork example provides informal motivation for the following two central points concerning the nature of confirmational independence and its intuitive relation to probabilistic screening-off:

¹⁰This theorem is closely related to a result reported in Eells and Fitelson (2000b) which says that each of our four Bayesian relevance measures — *except* r — satisfies the following *hypothesis symmetry* condition: (HS) $c(H, E | K) = -c(\bar{H}, E | K)$.

¹¹See Eells and Fitelson (2000b) for an independent set of high-level desiderata which also narrow the field to the two measures d and l . Pace Milne (1996), d and l seem, in many ways, to be the two most serious candidates for “the one true measure of confirmation.”

¹²Sober *presupposes* that the likelihood ratio (which is ordinally equivalent to l) is the correct way to measure degree of evidential support. I will use Sober’s example in what follows to motivate certain intuitive aspects of confirmational independence which will, ultimately, lead to an *argument in favor of* l (as opposed to d , r , or s) as a measure of evidential support.

- Confirmational independence is inherently a *three-place* relation. That is, when we say E_1 and E_2 are *confirmationally independent regarding H* , we are *not* saying that E_1 and E_2 are *unconditionally independent of each other*. We are talking about a kind of (ternary) independence that depends crucially on the hypothesis H .
- *Screening-off* of E_1 from E_2 by H (and by \bar{H}) is (intuitively) intimately connected with confirmational independence of E_1 and E_2 regarding H .

In the next section, I will describe a more general, probabilistic model that is intended to make the connection between probabilistic screening-off and confirmational independence more precise. This formal model will also allow us to generate concrete, numerical examples which will, ultimately, be used to show that only the log-likelihood ratio measure l properly handles the (general) relationship between probabilistic screening-off and confirmational independence.

2.3.2 A Formal Model

To formally motivate the general connection between probabilistic screening-off and confirmational independence, I will use a simple, abstract model. I will call this model the *urn model*.¹³ The background evidence for the urn model is assumed at the outset to contain the following information:

An urn has been selected at random from a collection of urns. Each urn contains some balls. In some of the urns the proportion of white balls to other balls is x and in all the other urns the proportion of white balls is y , $0 < x, y < 1$. The proportion of urns of the first type is z , $0 < z < 1$. Balls are to be drawn randomly from the selected urn, with replacement.

Let H be the hypothesis that the proportion of white balls in the urn is x . Let W_i state that the ball drawn on the i th draw ($i \geq 1$) is white. I take it as intuitively clear that W_1 and W_2 are mutually confirmationally independent regarding H , regardless of the values of x , y , and z .¹⁴ Hence, I propose the following adequacy condition for measures of degree of confirmation:

- (UC) If \mathfrak{c} is an adequate measure of degree of confirmation then, both $\mathfrak{c}(H, W_1 | W_2) = \mathfrak{c}(H, W_1)$, and $\mathfrak{c}(H, W_2 | W_1) = \mathfrak{c}(H, W_2)$ for all urn examples (regardless of the values of x , y , and z).

¹³The urn model is due to Patrick Maher.

¹⁴Ellery Eells (personal communication) worries that for extreme (or near extreme) values of x , y , or z , this intuition might break down. He may be right about this (although, as a defender of l , I will insist that any such breakdown can be explained away, and is probably just a psychological “edge effect”, owing to the extremity of the values of x , y or z , and not to considerations relevant to their confirmational independence *per se*). However, in the Appendix (Theorem 3), I show that the measures d , r , and s *fail* to obey this intuition, even in cases where the values of x , y , and z are all *far from* extreme. As a result, d , r , and s will not even judge E_1 and E_2 as confirmationally independent regarding H in *Sober’s* example. This seems highly unintuitive, and should cast serious doubt on the adequacy of d , r , and s .

What probabilistic feature of the urn model could be responsible for the (presumed) fact that W_1 and W_2 are confirmationally independent regarding H ? The feature cannot depend on the *values* of the probabilities involved, since we did not specify what these are except to say that they are not zero or one (a requirement imposed to ensure that the relevant conditional probabilities are all defined). Moreover, as we saw in Sober’s example, the feature cannot depend on the *unconditional* probabilistic independence of W_1 and W_2 , since W_1 and W_2 will *not*, in general, be independent of *each other* (e.g., if each of W_1 and W_2 *individually* confirms H). This does not leave much. Two considerations that remain are that the following two identities hold in all urn examples:

$$(1) \quad \Pr(W_1 \& W_2 \mid H) = \Pr(W_1 \mid H) \cdot \Pr(W_2 \mid H)$$

$$(2) \quad \Pr(W_1 \& W_2 \mid \bar{H}) = \Pr(W_1 \mid \bar{H}) \cdot \Pr(W_2 \mid \bar{H})$$

Identity (1) states that H *screens-off* W_1 from W_2 (or, equivalently, W_2 from W_1). Similarly, identity (2) states that \bar{H} likewise screens-off W_1 from W_2 . What I am suggesting, then, is that screening-off by H and \bar{H} is a *sufficient* condition for W_1 and W_2 to be mutually confirmationally independent regarding H . This suggests that (UC) might be strengthened to the following screening-off adequacy condition for measures of confirmation:

$$(SC) \text{ If } \mathfrak{c} \text{ is an adequate measure of confirmation, and if } H \text{ and } \bar{H} \text{ both screen-off } E_1 \text{ from } E_2, \text{ then } \mathfrak{c}(H, E_1 \mid E_2) = \mathfrak{c}(H, E_1) \text{ and } \mathfrak{c}(H, E_2 \mid E_1) = \mathfrak{c}(H, E_2).$$

I find (SC) an attractive principle; but, for the purposes of this paper, I will use only the weaker (and perhaps more intuitive) adequacy condition (UC).¹⁵ The following theorem states that the only measure among our four measures d , r , l , and s that satisfies (UC) is the log-likelihood ratio measure l .¹⁶

Theorem 3. *The measures d , r , and s violate (UC), but l satisfies (UC).*

Thus, only the log-likelihood ratio l satisfies the low-level, probabilistic screening-off desideratum. I think this is a compelling reason to favor the log-likelihood ratio measure over the other measures currently defended in the philosophical literature (at least, when it comes to judgments of confirmational independence

¹⁵Heckerman (1988, page 19) has suggested an adequacy condition that is equivalent to (SC). He gives no justification for this principle. I take the urn model to be a *partial* justification of (SC). However, I prefer the present approach since it makes use only of the weaker (and, I think, more intuitive) (UC). Incidentally, I do *not* think that screening-off on H and \bar{H} is a *necessary* condition for mutual confirmational independence regarding H (neither does Heckerman). I discuss this issue further in the Appendix, when I prove Theorem 3.

¹⁶Heckerman (1988) claims to prove a much more ambitious, and closely-related result. He claims to show that only measures that are ordinally equivalent to l satisfy (SC). Unfortunately, his argument is fallacious for subtle mathematical reasons — see Halpern (1996). In particular, Heckerman’s argument presupposes that an agent’s probability space is infinite, and satisfies some rather strong (unmotivated) mathematical constraints (Halpern 1996, pages 1318–1319). Unlike Heckerman’s argument, my argument makes use only of the finitistic adequacy condition (UC), and requires no additional, strong mathematical presuppositions.

regarding a hypothesis).¹⁷ As such, this provides a possible (at least, partial) solution to the problem of the plurality of Bayesian measures of confirmation described in Fitelson (1999). In the next section, I will discuss another application of my account of independent evidence.

3 An Application to Evidential Diversity

Philosophers of science dating back at least to Carnap (1945) have shared the intuition that collections of evidence that are ‘diverse’ or ‘varied’ should (*ceteris paribus*¹⁸) confirm more strongly than collections of evidence that are ‘narrow’ or ‘homogeneous’. I have elsewhere (see Fitelson 1996) called this the *confirmational significance of evidential diversity* (CSED). I suspect that the notion of *independent* evidence can undergird, at least partially, our intuitions about the significance of *diverse* evidence. At least one recent philosopher of science seems to share this suspicion. Sober (1989) shows (essentially¹⁹) that the log-likelihood ratio measure l satisfies the following condition:

- (\mathcal{D}) If each of E_1 and E_2 individually confirms H , and if E_1 and E_2 are confirmationally independent regarding H according to \mathfrak{c} , then $\mathfrak{c}(H \mid E_1 \& E_2) > \mathfrak{c}(H \mid E_1)$ and $\mathfrak{c}(H \mid E_1 \& E_2) > \mathfrak{c}(H \mid E_2)$.

It is a direct corollary of Theorem 1 that — according to *all three* measures of confirmation d , r , and l — two pieces of *independent* confirmatory evidence will always provide stronger confirmation than either one of them provides individually. In other words, we have already shown that the three most popular measures of confirmation d , r , and l *all* satisfy \mathcal{D} . It seems to me that \mathcal{D} could be used to provide a rather simple and elegant (partial²⁰) Bayesian account of

¹⁷The intimate connection between probabilistic screening-off of the kind described here and our intuitive judgments of independent inductive support has been pointed out by several recent authors (and used by some a reason to favor likelihood-ratio based measures of support), including: Good (1983), Pearl (1988), Heckerman (1988), and Schum (1994).

¹⁸See Fitelson (1996) for an elaboration of the *ceteris paribus* conditions that are tacitly presupposed in the Bayesian explication of CSED offered by Horwich (1982). I will later discuss the *ceteris paribus* clauses implicit in Howson and Urbach’s (1993) ‘correlation’ approach to CSED. Carnap’s original (1945, page 94) explication of CSED also requires some rather sophisticated *ceteris paribus* conditions. But, since Carnap’s original account of CSED does not make use of any of the measures d , r , l , or s , it is beyond the scope of this paper.

¹⁹Strictly speaking, Sober proves something *weaker* than this. He proves that l satisfies the consequent of \mathcal{D} under the *stronger* (wrt l) assumption that H (and \bar{H}) *screens-off* E_1 from E_2 . Our result is also more general than Sober’s in the sense that it applies not only to l but to d and r as well (*i.e.*, our result \mathcal{D} is *not* sensitive to the choice of measure of confirmation).

²⁰I do not mean to suggest that confirmational independence can be used to undergird *all* of our intuitions about the value of diverse evidence. But, I do think that there are many important scientific cases that fit this mold. For instance, the intuition that evidence from independent domains of application (*e.g.*, celestial *vs* terrestrial domains) of a theory often confirm more strongly than the same amount of evidence from domains of application that are not independent is a canonical example of the kind of intuition I have in mind here. Moreover, Sober (1989, page 124) explains how the notion of independent evidence regarding a hypothesis can be useful in the context of phylogenetic inference (*e.g.*, the problem of inferring the character states of ancestors from the observed character states of their descendants).

CSED. The basic idea behind such an approach would be that it is not evidence of different ‘kinds’ *per se* that will boost confirmational power. Rather, it is *data whose confirmational power is maximal, given the evidence we already have* that are confirmationally advantageous. And, \mathcal{D} provides a robust, general sufficient²¹ condition for this sort of confirmational boost.

It is *not* generally the case (as was pointed out by Carnap 1962) that two pieces of confirmatory evidence *simpliciter* will always provide stronger confirmation than just one. With \mathcal{D} , we have identified a very general sufficient condition for increased confirmational power. One nice feature of this sufficient condition is that *it does not depend sensitively on one’s choice of measure of confirmation*.²² Below, I compare the present approach to CSED with a recent Bayesian alternative proposed by Howson and Urbach (1993).

3.1 Comparison with the ‘correlation’ approach

Howson and Urbach (1993) propose a different way to account for our intuitions about CSED.²³ This approach asks us to consider *not* whether E_1 and E_2 are confirmationally independent *regarding H*. Rather, Howson and Urbach (1993) suggest that the important thing is whether or not E_1 and E_2 are *unconditionally stochastically independent*. Howson and Urbach (1993, pages 113–114, my italics) summarize the their ‘correlation’ account as follows:

Evidence that is varied is often regarded as offering better support to a hypothesis than an equally extensive volume of homogeneous evidence . . . According to the Bayesian, if two data sets are entailed by a hypothesis (or have similar probabilities relative to it²⁴), and one of them confirms more strongly than the other, this must be due to a corresponding difference between the data in their probabilities . . . The idea of similarity between items of evidence is expressed naturally in probabilistic terms by saying that e_1 and e_2 are similar if $P(e_2 | e_1)$ is higher than $P(e_2)$, and one might add that the more the first probability exceeds the second, the greater the similarity. *This means that e_2 would provide less support if e_1 had already been cited as evidence than if it was cited by itself.*

The most charitable interpretation of the above proposal of Howson and Urbach would seem to be the following rather complicated nested conditional:

²¹As was the case with (UC) and (SC), I am *not* claiming that \mathcal{D} is a *necessary* condition for increased confirmational power in this sense (indeed, it is *not* — proofs omitted).

²²In other words, the present explication of CSED is *robust* in the sense of Fitelson (1999). As we will see, the so-called “correlation” approach of Howson and Urbach (1993), is not robust in this sense, since their arguments tacitly presuppose the use of measure r .

²³Earman (1992) discusses a similar approach. Basically, the same criticisms will apply to his account. I will focus on the account of Howson and Urbach, since their characterization of the ‘correlation’ approach is closer in spirit to my presentation. See Forster (1995) for some good criticisms of Earman’s particular ‘correlation’ explication of CSED.

²⁴Howson and Urbach’s parenthetical remark that their ‘*ceteris paribus* condition’ can be weakened to $\Pr(E_1 | H) = \Pr(E_2 | H) = \Pr(E_1 \& E_2 | H)$ — while preserving the general truth of the main tenet of their account of CSED — is *false*. See footnote 36 in the Appendix (in the proof of Theorem 4) for a proof that this parenthetical remark is incorrect.

(\mathcal{H}) If the following probabilistic ‘*ceteris paribus* clause’ is satisfied:
 (CP) $\Pr(E_1 | H) = \Pr(E_2 | H) = \Pr(E_1 \& E_2 | H) = 1$,
 then if $\Pr(E_2 | E_1) > \Pr(E_2)$, then $\mathfrak{c}(H, E_2 | E_1) < \mathfrak{c}(H, E_2)$.²⁵

In other words, Howson and Urbach are claiming that (*ceteris paribus*)²⁶ pieces of evidence E_1 and E_2 that are *unconditionally* positively correlated will *not* be confirmationally independent regarding a hypothesis H (and, that E_1 and E_2 will tend to *cancel* each other’s support for H in such cases). I see several serious problems with Howson and Urbach’s proposal \mathcal{H} .²⁷

As we have already seen in Sober’s conjunctive fork example, pieces of *confirmationally independent* evidence will often be *unconditionally* positively correlated (and often *strongly* so). Newspaper reports (E_1) and radio reports (E_2) about the outcome (H) of a baseball game often fail to be *unconditionally* independent. This does nothing to undermine our intuition that E_1 and E_2 are *confirmationally independent* regarding H . Moreover, this example is representative of a wide range of cases. The conjunctive fork structure is common in (intuitive) examples of confirmational independence. For example, consider what doctors do when they seek independent confirmation of a diagnosis. They look for confirmationally independent corroborating symptoms. Such symptoms will typically be *unconditionally* correlated with already observed symptoms. But, *conditional on the relevant diagnostic hypothesis*, confirmationally independent symptoms will tend to be stochastically independent. It is *conditional* independence that is relevant here, not unconditional independence.

At best, Howson and Urbach have shown (*via* \mathcal{H}) that confirmational independence and unconditional stochastic dependence cannot co-occur in the extreme, deductive cases in which (CP) holds.²⁸ If \mathcal{H} were true for an interesting class of Bayesian confirmation measures \mathfrak{c} , then Howson and Urbach’s account would, at least, provide some useful information about the relationship between confirmational independence and unconditional stochastic independence in the case of deductive evidence. Unfortunately, as the following theorem states, among the four measures we have considered, Howson and Urbach’s \mathcal{H} is satisfied *only* by the log-ratio measure r , which we have already shown to be inadequate when it comes to judgments about confirmational independence.

Theorem 4. \mathcal{H} is true if $\mathfrak{c} = r$, but \mathcal{H} is false if $\mathfrak{c} = d$, $\mathfrak{c} = l$, or $\mathfrak{c} = s$.

²⁵In fact, Howson and Urbach seem to be making an even stronger, *quantitative* claim. They seem to be saying that if (CP) is satisfied, then the *greater* $\Pr(E_2 | E_1)$ is than $\Pr(E_2)$, the *lesser* $\mathfrak{c}(H, E_2 | E_1)$ will be than $\mathfrak{c}(H, E_2)$. I have chosen to criticize the (weaker) *qualitative* interpretation \mathcal{H} , since \mathcal{H} ’s falsity entails the falsity of the stronger, quantitative claim.

²⁶Howson and Urbach’s (CP) is just a bit stronger than the probabilistic *ceteris paribus* clause that is needed to shore-up Horwich’s (1982) account of CSED. See Fitelson (1996).

²⁷Note that Howson and Urbach’s \mathcal{H} only purports to explain why a *lack* of ‘diversity’ can be *bad*. \mathcal{H} cannot tell us why or how evidential ‘diversity’ can be *good*. In this sense, Howson and Urbach’s \mathcal{H} does not seem to directly address the traditional problem of CSED.

²⁸The fact that odd things can happen in such extreme cases was pointed out by Sober (1989, page 279). There, Sober explains that many of the salient epistemological differences between independent and dependent evidence collapse in the extreme (deterministic) case.

Howson and Urbach must either embrace the unattractive option of defending the measure r , or they must defend some other measure of confirmation which satisfies \mathcal{H} .²⁹ In either case, Howson and Urbach must reject the general connection (SC) between screening-off and confirmational independence, since (SC) and \mathcal{H} are logically incompatible in cases where both (CP) and screening-off obtain. That is, in the case of deterministic conjunctive forks, (SC) and \mathcal{H} cannot both be true.³⁰

4 Summary of Results

The following table summarizes the main results reported in this paper.

Name and Section of Condition \mathcal{C}	Is \mathcal{C} satisfied by the measure:			
	$d?$	$r?$	$l?$	$s?$
Peircean Additivity Condition \mathcal{A} (See §2.1 and Appendix §A for discussion)	YES	YES	YES	NO
Negation Symmetry Condition \mathcal{S} (See §2.2 and Appendix §B for discussion)	YES	NO	YES	YES
The Urn Condition (UC) (See §2.3.2 and Appendix §C for discussion)	NO	NO	YES	NO
Howson and Urbach's Condition \mathcal{H} (See §3.1 and Appendix §D for discussion)	NO	YES	NO	NO

5 Conclusion

I have outlined a general Bayesian account of confirmationally independent evidence regarding a hypothesis. At its heart, this account traces back to the pioneering work of C.S. Peirce. I have shown that a wide variety of (but, surprisingly, not all) Bayesian measures of degree of confirmation satisfy the most basic Peircean desiderata for adequate accounts of independent inductive support. I have also applied the idea of confirmational independence to two important problems in Bayesian confirmation theory: (*i*) the problem of the plurality of Bayesian measures of confirmation (as described in Fitelson 1999), and (*ii*) the problem of the confirmational significance of evidential diversity (as described in Fitelson 1996). I suspect that other useful applications of the present account of confirmational independence await discovery.

²⁹This would probably not be an easy task. Other than r (or measures ordinally equivalent to r), all relevance measures I have studied violate \mathcal{H} . See Appendix §D for more on \mathcal{H} .

³⁰This is easily proved. Assume that (CP) obtains (which implies that H screens-off E_1 from E_2), and that \bar{H} screens-off E_1 from E_2 . In such a case, H , E_1 , and E_2 will form a (deterministic) conjunctive fork. Now, if \mathcal{H} is true in such a case, then we must have $c(H, E_2 | E_1) < c(H, E_2)$. But, (SC) entails that in such a case $c(H, E_2 | E_1) = c(H, E_2)$. Therefore, in the case of deterministic conjunctive forks, (SC) and \mathcal{H} cannot both be true. \square

Appendix

A Proof of Theorem 1

Theorem 1. *Each of the measures d , r , and l satisfies \mathcal{A} , but s violates \mathcal{A} .*

Proof. This proof has four parts.³¹ The proofs for d and r are easy:

$$\begin{aligned}
 & d(H, E_1 | E_2) = d(H, E_1) \\
 \therefore \Pr(H | E_1 \& E_2) - \Pr(H | E_2) &= \Pr(H | E_1) - \Pr(H) \\
 (d) \quad \therefore \Pr(H | E_1 \& E_2) - \Pr(H) &= (\Pr(H | E_1) - \Pr(H)) \\
 & \quad + (\Pr(H | E_2) - \Pr(H)) \\
 \therefore d(H, E_1 \& E_2) &= d(H, E_1) + d(H, E_2)
 \end{aligned}$$

$$\begin{aligned}
 & r(H, E_1 | E_2) = r(H, E_1) \\
 \therefore \log[\Pr(H | E_1 \& E_2)] - \log[\Pr(H | E_2)] &= \log[\Pr(H | E_1)] - \log[\Pr(H)] \\
 (r) \quad \therefore \log[\Pr(H | E_1 \& E_2)] - \log[\Pr(H)] &= (\log[\Pr(H | E_1)] - \log[\Pr(H)]) \\
 & \quad + (\log[\Pr(H | E_2)] - \log[\Pr(H)]) \\
 \therefore r(H, E_1 \& E_2) &= r(H, E_1) + r(H, E_2)
 \end{aligned}$$

The proof for l is only slightly more involved. For the l case of the theorem, we will prove that the likelihood ratio (λ) is *multiplicative* under the assumption of confirmational independence. That the *logarithm* of λ (*i.e.*, l) is *additive* under the assumption of confirmational independence then follows straightaway.

$$\begin{aligned}
 & l(H, E_1 | E_2) = l(H, E_1) \\
 \therefore \lambda(H, E_1 | E_2) &= \lambda(H, E_1) \quad [\text{strict monotonicity of } \log(\bullet)] \\
 \therefore \frac{\Pr(E_1 | H \& E_2)}{\Pr(E_1 | \bar{H} \& E_2)} &= \frac{\Pr(E_1 | H)}{\Pr(E_1 | \bar{H})} \quad [\text{def. of } \lambda] \\
 (l) \quad \therefore \frac{\Pr(E_1 | H)}{\Pr(E_1 | \bar{H})} &= \frac{\Pr(E_1 \& E_2 | H)}{\Pr(E_1 \& E_2 | \bar{H})} \cdot \frac{\Pr(E_2 | \bar{H})}{\Pr(E_2 | H)} \quad [\text{def. of } \Pr(\bullet | \bullet)] \\
 \therefore \frac{\Pr(E_1 \& E_2 | H)}{\Pr(E_1 \& E_2 | \bar{H})} &= \frac{\Pr(E_1 | H)}{\Pr(E_1 | \bar{H})} \cdot \frac{\Pr(E_2 | H)}{\Pr(E_2 | \bar{H})} \\
 \therefore \lambda(H, E_1 \& E_2) &= \lambda(H, E_1) \cdot \lambda(H, E_2) \\
 \therefore l(H, E_1 \& E_2) &= l(H, E_1) + l(H, E_2) \quad [\text{additivity of } \log(\bullet)]
 \end{aligned}$$

The s case of the theorem is the trickiest, because it requires us to show that there is *no* (symmetric) monotonic function f such that, for all E_1 , E_2 , and H , if E_1 and E_2 are confirmationally independent regarding H according to s , then $s(H, E_1 \& E_2) = f[s(H, E_1), s(H, E_2)]$, where f is linear in some (monotonically) transformed space. Happily, I have proven the following *much stronger* result:

³¹Throughout the Appendix, we will suppress the contents of the background evidence K other than E_1 and E_2 . Moreover, we will try to prove the strongest results we know. Usually, these will be considerably stronger than the theorems that are stated in the main text.

(*) There exist probability models \mathcal{M}_1 and \mathcal{M}_2 such that:

\mathcal{M}_1	\mathcal{M}_2
$s(H, E_1 E_2) = s(H, E_1) = \frac{1}{4}$	$s(H, E_1 E_2) = s(H, E_1) = \frac{1}{4}$
$s(H, E_2 E_1) = s(H, E_2) = \frac{1}{4}$	$s(H, E_2 E_1) = s(H, E_2) = \frac{1}{4}$
$s(H, E_1 \& E_2) = \frac{15}{44} - \frac{96}{4451+3\cdot\sqrt{1254641}}$ ≈ 0.3286	$s(H, E_1 \& E_2) = \frac{15}{44} + \frac{96}{3\cdot\sqrt{1254641}-4451}$ ≈ 0.2529

Of course, it follows from (*) that there can be *no function f whatsoever* such that for all E_1, E_2 , and H , if E_1 and E_2 are confirmationally independent regarding H according to s , then $s(H, E_1 \& E_2) = f[s(H, E_1), s(H, E_2)]$. This is because (i) \mathcal{M}_1 and \mathcal{M}_2 are both such that E_1 and E_2 are confirmationally independent regarding H according to s , (ii) In \mathcal{M}_1 and \mathcal{M}_2 , $s(H, E_1)$ and $s(H, E_2)$ are *constant* at the same value of $\frac{1}{4}$, but (iii) The value of $s(H, E_1 \& E_2)$ in \mathcal{M}_1 is different from the value of $s(H, E_1 \& E_2)$ in \mathcal{M}_2 . So, whatever $s(H, E_1 \& E_2)$ is in cases where E_1 and E_2 are confirmationally independent regarding H according to s , it cannot (in general) be of the form $f[s(H, E_1), s(H, E_2)]$ for *any f whatsoever*, since functions cannot give different values for identical arguments. Due to space limitations, I will not display here all the calculations necessary to show that the models \mathcal{M}_1 and \mathcal{M}_2 reported below have the desired properties.³²

\mathcal{M}_1	
$\Pr(H \& \bar{E}_1 \& \bar{E}_2) = \frac{1}{100}$	$\Pr(H \& E_1 \& \bar{E}_2) = \frac{1}{1000}$
$\Pr(\bar{H} \& E_1 \& \bar{E}_2) = \frac{1}{100}$	$\Pr(H \& \bar{E}_1 \& E_2) = \frac{1}{1000}$
$\Pr(H \& E_1 \& E_2) = \frac{9\cdot(1183-\sqrt{1254641})}{70400}$	$\Pr(\bar{H} \& E_1 \& E_2) = \frac{87\cdot(1183-\sqrt{1254641})}{352000}$
$\Pr(\bar{H} \& \bar{E}_1 \& E_2) = \frac{1}{100}$	$\Pr(\bar{H} \& \bar{E}_1 \& \bar{E}_2) = \frac{121}{125} + \frac{3\cdot(\sqrt{1254641}-1183)}{8000}$

\mathcal{M}_2	
$\Pr(H \& \bar{E}_1 \& \bar{E}_2) = \frac{1}{100}$	$\Pr(H \& E_1 \& \bar{E}_2) = \frac{1}{1000}$
$\Pr(\bar{H} \& E_1 \& \bar{E}_2) = \frac{1}{100}$	$\Pr(H \& \bar{E}_1 \& E_2) = \frac{1}{1000}$
$\Pr(H \& E_1 \& E_2) = \frac{9\cdot(1183+\sqrt{1254641})}{70400}$	$\Pr(\bar{H} \& E_1 \& E_2) = \frac{87\cdot(1183+\sqrt{1254641})}{352000}$
$\Pr(\bar{H} \& \bar{E}_1 \& E_2) = \frac{1}{100}$	$\Pr(\bar{H} \& \bar{E}_1 \& \bar{E}_2) = \frac{121}{125} - \frac{3\cdot(\sqrt{1254641}+1183)}{8000}$

This completes the proof of Theorem 1. \square

B Proof of Theorem 2

Theorem 2. *Each of the measures d, l , and s satisfies \mathcal{S} , but r violates \mathcal{S} .*

³²The probability models in this Appendix were found and verified using *MATHEMATICA*.

Proof. This proof has four parts. The d , l , and s cases reduce to trivial algebraic identities (due to space limitations, I won't include here the easy proofs for these cases). For the r case, we need to show that there exists a probability model \mathcal{M} such that: (i) both $r(H, E_1 | E_2) = r(H, E_1)$ and $r(H, E_2 | E_1) = r(H, E_2)$, but (ii) either $r(\bar{H}, E_1 | E_2) \neq r(\bar{H}, E_1)$ or $r(\bar{H}, E_2 | E_1) \neq r(\bar{H}, E_2)$. Here is one such model \mathcal{M} (due to space limitations, I omit the computational details).

\mathcal{M}	
$\Pr(H \& \bar{E}_1 \& \bar{E}_2) = \frac{1}{64}$	$\Pr(H \& E_1 \& \bar{E}_2) = \frac{1}{64}$
$\Pr(\bar{H} \& E_1 \& \bar{E}_2) = \frac{1}{64}$	$\Pr(H \& \bar{E}_1 \& E_2) = \frac{87 + \sqrt{66265}}{704}$
$\Pr(H \& E_1 \& E_2) = \frac{1}{4}$	$\Pr(\bar{H} \& E_1 \& E_2) = \frac{1}{16}$
$\Pr(\bar{H} \& \bar{E}_1 \& E_2) = \frac{1}{8}$	$\Pr(\bar{H} \& \bar{E}_1 \& \bar{E}_2) = \frac{276 - \sqrt{66265}}{704}$

This completes the proof of Theorem 2. □

C Proof of Theorem 3

Theorem 3. *The measures d , r , and s violate (UC), but l satisfies (UC).*

Proof. For the d , r , and s cases of the theorem, it will suffice to produce an urn example (*i.e.*, an assignment of values on $(0, 1)$ to the variables x , y , and z) such that either $\mathfrak{c}(H, W_1 | W_2) \neq \mathfrak{c}(H, W_1)$ or $\mathfrak{c}(H, W_2 | W_1) \neq \mathfrak{c}(H, W_2)$, for each of the three measures d , r , and s . The following (far from extreme³³) assignment does the trick: $\langle x, y, z \rangle = \langle \frac{1}{2}, \frac{49}{100}, \frac{1}{2} \rangle$. On this assignment, we have the following salient probabilistic facts (computational details omitted for reasons of space):

- (d) $d(H, W_2 | W_1) = 2450/485199 < d(H, W_2) = 1/198$
- (r) $r(H, W_2 | W_1) = \log(4950/4901) < r(H, W_2) = \log(100/99)$
- (s) $s(H, W_2 | W_1) = 245000/24500099 < s(H, W_2) = 100/9999$

For the l case, we will show that l satisfies the stronger condition (SC).³⁴

$$\begin{aligned}
 & \Pr(E_1 | H \& E_2) = \Pr(E_1 | H) \quad [\text{screening-off assumption}] \\
 & \Pr(E_1 | \bar{H} \& E_2) = \Pr(E_1 | \bar{H}) \quad [\text{screening-off assumption}] \\
 (l) \quad & \therefore \frac{\Pr(E_1 | H \& E_2)}{\Pr(E_1 | \bar{H} \& E_2)} = \frac{\Pr(E_1 | H)}{\Pr(E_1 | \bar{H})} \\
 & \therefore l(H, E_1 | E_2) = l(H, E_1)
 \end{aligned}$$

It is easy to show that, for any of the three measures d , r , or l (but *not* for s), $\mathfrak{c}(H, E_1 | E_2) = \mathfrak{c}(H, E_1)$ iff $\mathfrak{c}(H, E_2 | E_1) = \mathfrak{c}(H, E_2)$. That, together with the reasoning above, completes the l case, and with it the proof of Theorem 3. □

³³ y can be *arbitrarily close* to $\frac{1}{2}$, while preserving the counterexample. See footnote 14.

³⁴It is interesting to note that l does *not* satisfy the converse of (SC), or the converse of (UC). This is why I do *not* take screening-off to be *necessary* for confirmational independence.

D Proof of Theorem 4

Theorem 4. \mathcal{H} is true if $\mathbf{c} = r$, but \mathcal{H} is false if $\mathbf{c} = d$, $\mathbf{c} = l$, or $\mathbf{c} = s$.³⁵

Proof. For the r case of the theorem, we begin by assuming that the probabilistic ‘*ceteris paribus* clause’ (CP) is satisfied. That is, we assume: $\Pr(E_1 | H) = \Pr(E_2 | H) = \Pr(E_1 \& E_2 | H) = 1$. Then, we apply (CP), the definition of r , and Bayes’ Theorem to derive the following pair of probabilistic facts:

$$\begin{aligned}
 (3) \quad r(H, E_2 | E_1) &= \log \left[\frac{\Pr(H | E_1 \& E_2)}{\Pr(H | E_1)} \right] \\
 &= \log \left[\frac{\Pr(E_1 \& E_2 | H) \cdot \Pr(H) \cdot \Pr(E_1)}{\Pr(E_1 \& E_2) \cdot \Pr(E_1 | H) \cdot \Pr(H)} \right] \\
 &= \log \left[\frac{\Pr(E_1)}{\Pr(E_1 \& E_2)} \right] \\
 &= \log \left[\frac{1}{\Pr(E_2 | E_1)} \right]
 \end{aligned}$$

$$\begin{aligned}
 (4) \quad r(H, E_2) &= \log \left[\frac{\Pr(H | E_2)}{\Pr(H)} \right] \\
 &= \log \left[\frac{\Pr(E_2 | H) \cdot \Pr(H)}{\Pr(E_2) \cdot \Pr(H)} \right] \\
 &= \log \left[\frac{1}{\Pr(E_2)} \right]
 \end{aligned}$$

Finally, we assume that E_1 and E_2 are positively correlated under \Pr . Or, more formally, we assume that $\Pr(E_2 | E_1) > \Pr(E_2)$. In conjunction with facts (3) and (4) above, this yields $r(H, E_2 | E_1) < r(H, E_2)$, as desired.³⁶

For the d , l , and s cases, it will suffice to produce a probability model in which (i) $\Pr(E_1 | H) = \Pr(E_2 | H) = \Pr(E_1 \& E_2 | H) = 1$, (ii) $\Pr(E_2 | E_1) > \Pr(E_2)$, but (iii) $\mathbf{c}(H, E_2 | E_1) \geq \mathbf{c}(H, E_2)$, for $\mathbf{c} = d$, $\mathbf{c} = l$, and $\mathbf{c} = s$. The following example does the trick. A card is drawn at random from a standard deck. Let H be the hypothesis that the card is the $\text{Q}\spadesuit$, E_1 be the proposition that the card is either a 10 or a face card, and E_2 be the proposition that the card is either a \heartsuit , or the $\text{Q}\spadesuit$, or the $\text{9}\spadesuit$. Due to space limitations, I omit the calculations which show that this example has the desired properties (i)–(iii) listed above.

This completes the proof of Theorem 4, as well as the Appendix. \square

³⁵ \mathcal{H} is also false for Carnap’s (1962, §67) relevance measure \mathbf{r} (as the example below shows).

³⁶ Notice that Howson and Urbach’s claim that (CP) can be weakened even further to (CP') $\Pr(E_1 | H) = \Pr(E_2 | H) = \Pr(E_1 \& E_2 | H)$ — while still preserving the truth of the $\mathbf{c} = r$ case of Theorem 4 — is *false*. If we only assume (CP'), then we will need to establish that $\Pr(E_2 | E_1) > \frac{\Pr(E_2)}{\Pr(E_2 | H)}$, in order to prove that $r(H, E_2 | E_1) < r(H, E_2)$. Unfortunately, $\Pr(E_2 | E_1) > \frac{\Pr(E_2)}{\Pr(E_2 | H)}$ does *not* follow from the fact that E_1 and E_2 are positively correlated under \Pr , unless one also assumes that $\Pr(E_2 | H) = 1$, which brings us back to (CP).

References

- Carnap, R. (1945). On inductive logic. *Philosophy of Science* 12, 72–97.
- Carnap, R. (1962). *Logical Foundations of Probability* (Second ed.). Chicago: University of Chicago Press.
- Christensen, D. (1999). Measuring confirmation. *Journal of Philosophy* XCVI, 437–61.
- Earman, J. (1992). *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory*. Cambridge: MIT Press.
- Eells, E. (1982). *Rational Decision and Causality*. Cambridge: Cambridge University Press.
- Eells, E. and B. Fitelson (2000a). Measuring confirmation and evidence. Unpublished manuscript.
- Eells, E. and B. Fitelson (2000b). Symmetries and asymmetries in evidential support. Unpublished manuscript.
- Fitelson, B. (1996). Wayne, Horwich, and evidential diversity. *Philosophy of Science* 63, 652–660.
- Fitelson, B. (1999). The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science* 66 (Proceedings), S362–S378.
- Forster, M. (1995). Bayes and bust: Simplicity as a problem for a probabilist’s approach to confirmation. *The British Journal for the Philosophy of Science* 46, 399–424.
- Gillies, D. (1986). In defense of the Popper-Miller argument. *Philosophy of Science* 53, 110–113.
- Good, I. (1983). *Good Thinking*. Minneapolis: University of Minnesota Press.
- Good, I. (1984). The best explicatum for weight of evidence. *Journal of Statistical Computation and Simulation* 19, 294–299.
- Halpern, J. (1996). A counterexample to theorems of Cox and Fine. In *AAAI-96*, pp. 1313–1319. Cambridge, Mass.: MIT Press.
- Heckerman, D. (1988). An axiomatic framework for belief updates. In L. Kanal and J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence 2*, pp. 11–22. New York: Elsevier Science Publishers.
- Horvitz, E. and D. Heckerman (1986). The inconsistent use of certainty measures in artificial intelligence research. In L. Kanal and J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence 1*, pp. 137–151. New York: Elsevier Science Publishers.

- Horwich, P. (1982). *Probability and Evidence*. Cambridge: Cambridge University Press.
- Howson, C. and P. Urbach (1993). *Scientific Reasoning: The Bayesian Approach*. La Salle: Open Court.
- Jeffrey, R. (1992). *Probability and the Art of Judgment*. Cambridge: Cambridge University Press.
- Joyce, J. (1999). *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Kemeny, J. and P. Oppenheim (1952). Degrees of factual support. *Philosophy of Science* 19, 307–324.
- Keynes, J. (1921). *A Treatise on Probability*. London: Macmillan.
- Kyburg, H. (1983). Recent work in inductive logic. In T. Machan and K. Lucey (Eds.), *Recent Work in Philosophy*, pp. 87–150. Lanham: Rowman & Allanheld.
- Mackie, J. (1969). The relevance criterion of confirmation. *The British Journal for the Philosophy of Science* 20, 27–40.
- Milne, P. (1996). $\log[p(h/eb)/p(h/b)]$ is the one true measure of confirmation. *Philosophy of Science* 63, 21–26.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco: Morgan Kaufman.
- Peirce, C. (1878). The probability of induction. *Popular Science Monthly* 12, 705–718.
- Pollard, S. (1999). Milne’s measure of confirmation. *Analysis* 59, 335–337.
- Reichenbach, H. (1956). *The Direction of Time*. University of California.
- Rosenkrantz, R. (1994). Bayesian confirmation: Paradise regained. *The British Journal for the Philosophy of Science* 45, 467–476.
- Schlesinger, G. (1995). Measuring degrees of confirmation. *Analysis* 55, 208–212.
- Schum, D. (1994). *The Evidential Foundations of Probabilistic Reasoning*. New York: John Wiley & Sons.
- Sober, E. (1989). Independent evidence about a common cause. *Philosophy of Science* 56, 275–287.
- Wayne, A. (1995). Bayesianism and diverse evidence. *Philosophy of Science* 62, 111–121.