

STUDIES IN BAYESIAN CONFIRMATION THEORY

BRANDEN FITELSON

Department of Philosophy
San José State University

branden@fitelson.org
http://fitelson.org/

Some Bayesian Background I

- Orthodox Bayesianism (*i.e.*, Bayesian *epistemology*) assumes that the degrees of belief (or credence) of rational agents are (*Kolmogorov* [29]) *probabilities*.
- $\Pr_a(H | K)$ denotes an (rational) agent *a*'s degree of credence in *H*, given the corpus *K* of background knowledge/evidence (called *a*'s "prior" for *H*).
- $\Pr_a(H | E \& K)$ denotes *a*'s degree of credence in *H* (relative to *K*) given that (or on *a*'s supposition that) *E*. This is also the agent's degree of belief in *H* (relative to *K*) upon learning *E* (called *a*'s "posterior" for *H*, on *E*, given *K*).
 - * Credences are *Kolmogorov* [21], [27], [12], *probabilities* [50], on *crisp sets* [54].
 - * Agents learn (*with certainty* [26]) via *conditionalization* [32], [33].
 - * "Priors" (\therefore *Bayesianism itself*) are *subjective* [49], [47], [34], [30], [6].
- I will bracket all of these issues. The problem I'm discussing only gets *worse* if Bayesianism is made more sophisticated along any of these dimensions!
- For simplicity, I will assume there is a *single* rational Bayesian probability function \Pr (and I'll drop the subscript "*a*" and the background corpus "*K*").

Some Bayesian Background II

- In (contemporary) Bayesian confirmation theory, evidence *E* *confirms* (or *supports*) a hypothesis *H* if learning *E* *raises the probability of H*.
- If learning *E* *lowers* the probability of *H*, then *E* *disconfirms* (or *counter-supports*) *H*, and if learning *E* *does not change* the probability of *H*, then *E* is confirmationally *neutral* regarding *H*. This is a *Pr-relevance* theory.
- Within (*Kolmogorov!* [10], [12]) probability theory, there are many *logically equivalent* ways of saying that *E* confirms *H*. Here are a few:
 - * *E* confirms *H* if $\Pr(H | E) > \Pr(H)$.
 - * *E* confirms *H* if $\Pr(E | H) > \Pr(E | \neg H)$.
 - * *E* confirms *H* if $\Pr(H | E) > \Pr(H | \neg E)$.
- By taking differences, (log-)ratios, *etc.*, of the left/right sides of these (or other equiv.) inequalities, a *plethora* of candidate *relevance measures of degree of confirmation* can be formed. (\mathcal{R}) $c(H, E) \lesseqgtr 0$ if $\Pr(H | E) \lesseqgtr \Pr(H)$.

Four Popular and Representative Relevance Measures

- *Dozens* of Bayesian relevance measures have been proposed in the philosophical literature (see [31] for a survey). Here are four popular ones.^a
 - * *Difference*: $d(H, E) =_{df} \Pr(H | E) - \Pr(H)$
 - * *Log-Ratio*: $r(H, E) =_{df} \log \left[\frac{\Pr(H | E)}{\Pr(H)} \right]$
 - * *Log-Likelihood-Ratio*: $l(H, E) =_{df} \log \left[\frac{\Pr(E | H)}{\Pr(E | \neg H)} \right]$
 - * "Normalized Difference" *d*: $s(H, E) =_{df} \Pr(H | E) - \Pr(H | \neg E) = \frac{1}{\Pr(\neg E)} \cdot d(H, E)$
- Logs are taken to ensure easy satisfaction of relevance criterion (\mathcal{R}). They are merely a useful convention (they're inessential, but they simplify things).
- The first part of our story concerns the *disagreement* exhibited by these measures, and its ramifications for Bayesian confirmation theory ...

^aUsers of *d* include [7], [6], and [26]. Users of *r* include [24], [35], and [25]. Users of *l* include [19], [48], and [15]. Users of *s* include [27] and [5]. See [10], [13], and [15] for further references.

Disagreement Between Alternative Relevance Measures

- What kind of disagreement between relevance measures is important?
- Mere *numerical* (or *conventional* or *syntactical*) differences between measures are not important, since they need not effect *ordinal* judgments of what is more/less well confirmed than what (by what).
- *Ordinal* differences are crucial, since they can effect the cogency of many arguments surrounding Bayesian confirmation theory.
- For instance, it is part of Bayesian lore that the observation of a black raven (E_1) confirms the hypothesis (H) that all ravens are black *more strongly than* the observation of a white shoe (E_2) does (given “actual corpus” K).
- But, given the standard background assumptions (K) in Bayesian accounts of Hempel’s ravens paradox, this conclusion [$c(H, E_1) > c(H, E_2)$] follows only for *some* measures of confirmation c (and *not* others).
- Such arguments are said to be *sensitive to choice of measure* [13].

The Problem of Measure Sensitivity

- A detailed study of the literature shows that *virtually every argument* involving quantitative Bayesian confirmation theory is sensitive to choice of measure [13]! Here are my favorites (I’ll briefly discuss three of these):
 - * Horwich [24] *et al.* on Hempel’s Ravens Paradox
 - * Horwich [24] *et al.* on the Confirmational Value of Varied Evidence
 - * The Popper-Miller Argument *Against* Bayesianism [41], [16]
 - * Rosenkrantz [46] and Earman [6] on the Problem of “Irrelevant Conjunction”
 - * Eells [7] and Sober [52] on Goodman’s “Grue” Paradox
 - * Earman [6] on the problem of (quantitative) old evidence
- There are many other important measure-sensitive arguments [3], [1], [2].
- One needn’t gerrymander or comb the historical literature for Bayesian relevance measures which fail to undergird these arguments.
- Each of these arguments is valid with respect to *only some* of d , r , l , and s .

Horwich *et al.* on Ravens & Variety of Evidence

- Almost all Bayesian accounts of both the Ravens Paradox and the value of “varied” evidence (*i.e.*, why more “varied” evidence E_1 is more confirmationally powerful than less “varied” evidence E_2) presuppose:

$$(1) \quad \text{If } \Pr(H | E_1) > \Pr(H | E_2), \text{ then } c(H, E_1) > c(H, E_2).$$

- The “normalized” difference measure s *violates* (1).^a
- Typically, the advocates of such arguments have used either d or r in their arguments (as it turns out, d , r , and l *all* satisfy (1)).
- None of these authors seems to provide (*independent*) reasons to prefer their measures over s (or other measures which violate (1)).
- In my [14] and [15], I propose a novel Bayesian explication of the confirmational value of *independent* evidence, based on l .

^aSo do Carnap’s [4, §67] relevance measure $r(H, E) = \Pr(H \& E) - \Pr(H) \cdot \Pr(E)$, Mortimer’s [36] measure $\Pr(E | H) - \Pr(E)$, and Nozick’s [37] measure $\Pr(E | H) - \Pr(E | \neg H)$.

The Popper-Miller Argument *Against* Bayesianism

- It isn’t just arguments/accounts *within* Bayesian confirmation theory that are sensitive to choice of measure. Some well-known *criticisms* of Bayesianism also rest on measure sensitive arguments.
- Most famously, Popper and Miller ([41], [16]) use the following property of the difference measure d to argue *against* Bayesianism (generally):

$$(2) \quad d(H, E) = d(H \vee E, E) + d(H \vee \neg E, E).$$
- As it turns out, neither the log-ratio measure r [42], nor the log-likelihood-ratio measure l [20] satisfies property (2).
- \therefore The Popper-Miller argument is *sensitive to choice of measure*.
- In the absence of reasons to think that d is a more accurate (and charitable) reconstruction of Bayesian confirmation theory than either r or l , the Popper-Miller argument remains (at best) *enthymematic*.

Tabular Summary of Some Measure-Sensitive Arguments

Argument	Valid wrt relevance measure:			
	$d?$	$r?$	$l?$	$s?$
Rosenkrantz on Irrelevant Conjunction	YES	NO	NO	YES
Earman on Irrelevant Conjunction	YES	NO	YES	YES
Eells on the Grue Paradox	YES	NO	NO	YES
Sober on the Grue Paradox	YES	NO	YES	YES
Horwich <i>et al.</i> on Ravens & Variety	YES	YES	YES	NO
Popper-Miller's Critique of Bayesianism	YES	NO	NO	YES
Earman's Old Evidence Critique of Bayesianism	YES	YES	NO	NO

Some Attempts to Resolve the Measure-Sensitivity Problem

- There do exist a few general arguments in the literature which aim to rule-out all but a small class of ordinally equivalent measures (*e.g.*, Milne [35], Good [18], Carnap [4], Kemeny & Oppenheim [28], and Heckerman [22]).
- Others have given “piecemeal” arguments which attack a *particular* class of measures, but fail to rule-out other competing measures (*e.g.*, Rosenkrantz [46], Earman [6], Gillies [16], Eells, and Sober [52]).
- In my dissertation [15], I provide a thorough survey of both kinds of arguments, and I show that none of them is completely satisfactory.
- Most notably, I have seen (in the literature^a) *no* compelling reasons to prefer the difference measure d over either l or s .
- Until such reasons are provided, the arguments of Gillies, Rosenkrantz, Eells, Horwich *et al.* will remain *enthymematic*.

^aRecent joint work of Eells & Fitelson in [8] and [9] has filled this gap in the literature.

Narrowing The Field I: Symmetries and Asymmetries in Evidential Support^a

- Consider the following two propositions concerning a card c , drawn at random from a standard deck of playing cards (classical model \mathcal{M}):
 E : c is the ace of spades. H : c is *some* spade.
- I take it as intuitively clear and uncontroversial that:
 1. The degree to which E supports $H \neq$ the degree to which H supports E , since $E \vDash H$, but $H \not\vDash E$. Intuitively, we have $c(H, E) \gg c(E, H)$.
 2. The degree to which E confirms $H \neq$ the degree to which $\neg E$ disconfirms H , since $E \vDash H$, but $\neg E \not\vDash \neg H$. Intuitively, $c(H, E) \gg -c(H, \neg E)$.
- Therefore, *no adequate relevance measure of support c should be such that either $c(H, E) = c(E, H)$ or $c(H, E) = -c(H, \neg E)$ (for all E and H and K).*
- Note: for all H, E (and K), $r(H, E) = r(E, H)$ and $s(H, E) = -s(H, \neg E)$. Both d and l satisfy both of these (a)symmetry desiderata.

^aThis slide is drawn from recent joint work of Eells & Fitelson [9].

Narrowing The Field II: Our Relevance Measures as Generalizations of Entailment

- $l(H, E) = \begin{cases} +\infty & \text{if } E \vDash H, \Pr(E) > 0, \Pr(H) \in (0, 1) \\ 0 & \text{if } E \perp\!\!\!\perp H, \Pr(E) > 0, \Pr(H) \in (0, 1) \\ -\infty & \text{if } E \vDash \neg H, \Pr(E) > 0, \Pr(H) \in (0, 1) \end{cases}$
- $d(H, E) = \begin{cases} \Pr(\neg H) & \text{if } E \vDash H, \Pr(E) > 0 \\ 0 & \text{if } E \perp\!\!\!\perp H, \Pr(E) > 0 \\ -\Pr(H) & \text{if } E \vDash \neg H, \Pr(E) > 0 \end{cases}$
- $r(H, E) = \begin{cases} \log \left[\frac{1}{\Pr(H)} \right] & \text{if } E \vDash H, \Pr(E) > 0, \Pr(H) > 0 \\ 0 & \text{if } E \perp\!\!\!\perp H, \Pr(E) > 0, \Pr(H) > 0 \\ -\infty & \text{if } E \vDash \neg H, \Pr(E) > 0, \Pr(H) > 0 \end{cases}$
- $s(H, E) = \begin{cases} \Pr(\neg H | \neg E) & \text{if } E \vDash H, \Pr(E) \in (0, 1) \\ 0 & \text{if } E \perp\!\!\!\perp H, \Pr(E) \in (0, 1) \\ -\Pr(H | \neg E) & \text{if } E \vDash \neg H, \Pr(E) \in (0, 1) \end{cases}$


Narrowing The Field III: Independent Evidence I

- Wittgenstein [53] alludes to a man who is doubtful about the reliability of a story he reads in the newspaper, so he buys another copy of the same issue of the same newspaper to double check.
- To fix our ideas, let's assume that the story in the NYT reports that (H) the Yankees won the world series. Let E_n be the evidence obtained by reading the n^{th} copy of the same issue of the NYT.
- Intuitively, the degree to which the conjunction $E_1 \& E_2$ confirms H is no greater than the degree to which E_1 alone confirms H .
- Also, it seems intuitive that an *independent* report E' (say, one heard on a NPR broadcast) *would* corroborate the NYT story.
- So, it seems intuitive that the degree to which $E_1 \& E'$ confirms H is greater than the degree to which E_1 alone confirms H .

Narrowing The Field III: Independent Evidence II

- How can we explain the epistemic difference between these two examples? Intuitively, a NYT report (E) and a NPR report (E') are *independent* in a way that two NYT reports (E_1, E_2) are not.
- It is *not* that the NYT report and the NPR report are Pr-independent *unconditionally*, since (far more often than not) the two reports will tend to agree. So, what kind of independence is at work here?
- As Sober [51] explains, the relevant probabilistic fact is that E and E' are independent — *given H* (or $\neg H$). That is, if we know the truth-value of H , then the dependence (correlation) between E and E' disappears.
- H *explains* the correlation between E and E' . This is because E and E' are *joint effects of the common cause H* . As Reichenbach [43] taught us: *common causes screen-off their joint effects from each other*.


Narrowing The Field III: Independent Evidence III




- E_1 and E_2 are *dependent* — *even if we know the truth-value of H* (perhaps if we knew the state of the NYT printing press just prior to publication, then *this* would render E_1 and E_2 independent).
- But, E_1 and E' are *independent* — once we know the truth-value of H . When this happens, we say that H *screens-off* E_1 from E' or that E_1 and E' are *conditionally independent, given H* (or $\neg H$).
- This kind of structure is a “Bayesian Network” [38], [22], [23], [2].

Narrowing The Field III: Independent Evidence IV

- If two pieces of (confirmatory) evidence E_1 and E_2 are independent regarding a hypothesis H , then the conjunction $E_1 \& E_2$ should confirm H more strongly than either conjunct does severally:



- More precisely, (as C.S. Peirce [39] suggests) the degree of support provided by the conjunction $E_1 \& E_2$ should simply be the *sum* of the several degrees of support provided by each conjunct:



Narrowing The Field III: Independent Evidence V

- Measures d , r , and l satisfy the first of these desiderata (s does *not*!) [15]. Only measure l satisfies the second (additivity) desideratum [14], [22].
- This provides a novel way of adjudicating between d and l (indeed, this adjudicates between l and anything else, for that matter).
- These ideas about independent evidence can also serve to ground a novel (partial) explication of the confirmational value of “varied” evidence.
- If “varied” evidence are *independent*, then they will provide a stronger confirmational boost than “narrow” or *dependent* evidence will provide — as measured by the log-likelihood-ratio l [14].
- According to l (but *not* according to d), strong independent confirmational boosts can be provided — even to hypotheses with high priors (e.g., Newton’s gravitational theory H , planetary data E_1 , and cometary data E_2 , in the year 1758 when Halley’s comet returned).

Extra Slide #1: Degree of *Belief* vs Degree of *Support*

- “But wait a minute! l has *unintuitive* behavior. If $\Pr(H)$ is very high, then H ’s probability *cannot* be raised very much, and so (*intuitively*) H *cannot* be confirmed very strongly. But, l *denies* this (whereas d gets this ‘right’!)”
- If the question is “How strongly should we *believe* H , given the supposition that E is true,” then the answer will be “very strongly” if $\Pr(H)$ is high (and E does not disconfirm H). This is *not* the issue.
 - I’m asking “How strong is the *relation of support* between E and H ” (assuming E does not disconfirm H). And, this does *not* always depend on the prior probability of H — contrary to what d presupposes.
 - When $E \vDash H$ (*relevantly!*), the degree to which E supports H is *maximally strong*, and does *not* depend on $\Pr(H)$. But, $d(H, E) = 1 - \Pr(H)$ here.
 - Also, I think that in the case of *independent* evidence, the degree of support does *not* depend on $\Pr(H)$. Moreover, if we talk about *odds* rather than \Pr , this “edge effect” disappears (it is an *artifact* of $[0,1]$ \Pr -scale).

Extra Slide #2: Some Details on Our l and the Limiting (Deductive) Case

- I claimed that, when $E \vDash H$, $l(H, E)$ is maximal ($+\infty$), and does not depend on the prior probability of H [if $\Pr(E) > 0$, $\Pr(H) \in (0, 1)$].
- But, if $E \vDash H$, then $\Pr(E | \neg H) = 0$. Shouldn’t we say that $l(H, E)$ is *undefined* in such cases, since it has a zero denominator?
- There are two ways to handle this. First, one could maintain that, *in the limit as $\Pr(E \& \neg H)$ approaches zero*, $l(H, E)$ diverges ($+\infty$). So, $l(H, E)$ is maximal and doesn’t depend on $\Pr(H)$ in such cases.
- Or, more satisfyingly, one could use the alternative measure [28]:

$$l'(H, E) = \frac{\Pr(E | H) - \Pr(E | \neg H)}{\Pr(E | H) + \Pr(E | \neg H)}$$

It is easy to show that (i) l' is *ordinally equivalent* to^a l , and (ii) l' takes on the values $+1/-1$ in cases where E (*relevantly*) entails/refutes H .

^aThis is because l' is a monotone increasing function of l [viz., $l' = \tanh(l/2)$]; see [19].

Extra Slide #3: Mere Syntactical Differences Between Measures

- Most existing Bayesian measures of support (*not s !*) can be written as some combination of (arithmetic) functions of the posterior and prior probabilities of H [*i.e.*, some function of $\Pr(H | E)$ and $\Pr(H)$].
- For instance, the likelihood ratio can be written as follows:

$$\frac{\Pr(E | H)}{\Pr(E | \neg H)} = \frac{\Pr(H | E) \cdot [1 - \Pr(H)]}{[1 - \Pr(H | E)] \cdot \Pr(H)}$$

- One should not put too much weight on *mere syntactical differences* between measures. It’s their *ordinal structure* that matters . . .
- For instance, one should not conclude that a measure doesn’t *depend* on $\Pr(H)$, simply because *one syntactical formulation* of the measure doesn’t happen to contain the string of symbols “ $\Pr(H)$ ”.

$$\frac{\Pr(H | E)}{\Pr(H)} = \frac{\Pr(E | H)}{\Pr(E)}$$

Extra Slide #4: The Problem of “Irrelevant Conjunction”

- According to deductive accounts of confirmation (e.g., Hempel’s H-D account), E confirms H (roughly) iff $H \vDash E$.
- Such accounts of confirmation have the following consequence:

(3) If E confirms H , then E confirms $H \& X$, for any X .
- While (3) is *not* a consequence of Bayesian confirmation, the following is:

(4) If $H \vDash E$, then E confirms $H \& X$, for any X .
- Bayesians try to mitigate the effects of (4), by arguing that:

(5) If $H \vDash E$, then $c(H \& X, E) < c(H, E)$, for any X .
- Such arguments have two problems: (i) they are sensitive to choice of measure c , and (ii) they make no appeal to the *irrelevance* of X .
- In [11], I give a new-and-improved Bayesian account.

Extra Slide #5: Rosenkrantz on “Irrelevant Conjunction”

- Rosenkrantz [46] provides a Bayesian resolution of the problem of Irrelevant Conjunction (*a.k.a.*, the Tacking Problem) which trades on the following property of the difference measure:

(6) If $H \vDash E$, then $d(H \& X, E) = \Pr(X | H) \cdot d(H, E)$.
- Neither r nor l satisfies property (6) [11].
- Rosenkrantz does provide some (pretty good) reasons to reject r . However, he [45] explicitly admits that he knows of “no compelling considerations that adjudicate between” d and l .
- So, it is (at best) unclear how one might consistently complete Rosenkrantz’s enthymematic treatment of the tacking problem.
- What’s worse, as I will explain later, I think there are good reasons to favor l over d as a measure of support.

Extra Slide #6: Earman on “Irrelevant Conjunction”

- Earman [6] gives a more robust resolution of the tacking problem which requires only the following logically weaker cousin of (6):

(6′) If $H \vDash E$, then $d(H \& X, E) < d(H, E)$.
- r violates even this weaker condition, but l satisfies (6′) [11].
- In this sense, Earman’s account is *less* sensitive to choice of measure (*i.e.*, more robust) than Rosenkrantz’s is.
- Earman’s account can be bolstered by providing compelling independent reasons to favor d (or l) over r (e.g., see below).
- Unfortunately, even the bolstered version of Earman’s account is inadequate. I provide a new and improved Bayesian resolution of the problem of irrelevant conjunction in [11].

Extra Slide #7: Goodman’s “Grue” Paradox

- Goodman presents an example involving the following two hypotheses (H and H') and observation report (E):

H : All emeralds are green.
 H' : All emeralds are grue.
 E : All emeralds that have been observed are green (\therefore grue).
- Where, the predicate “grue” is defined as follows:

x is grue if and only if either (i) x has been observed and x is green, or (ii) x has not been observed and x is blue.
- Bayesian answers to Goodman’s “new riddle of induction” have aimed to establish that H is better supported by E than H' is. That is, Bayesians have tried to show that $c(H, E) > c(H', E)$.
- As we have seen, at least two Bayesian accounts along these lines (those of Eells and Sober) are sensitive to choice of measure c .

Extra Slide #8: Eells on Goodman's "Grue" Paradox

- Eells [7] offers a Bayesian account of the Grue paradox (*a.k.a.*, Goodman's "new riddle of induction") which trades on the following property of the difference measure [where β, δ are:

$\beta =_{df} \Pr(H_1 \& E) - \Pr(H_2 \& E)$, and $\delta =_{df} \Pr(H_1 \& \neg E) - \Pr(H_2 \& \neg E)$]:

$$(7) \quad \text{If } \beta > \delta \text{ and } \Pr(E) < \frac{1}{2}, \text{ then } d(H_1, E) > d(H_2, E).$$

- Neither r nor l satisfies property (7).
- Eells does provide reasons (as reported in a paper by Sober, see below) to prefer the difference measure d over the log-ratio measure r , but he does not supply reasons to prefer d over l .
- Pending such reasons, Eells's argument remains *enthymematic*.
- Moreover, I will later provide reasons to favor l over d .

Extra Slide #9: Sober on Goodman's "Grue" Paradox

- Sober [52] describes a more robust Bayesian account of the Grue paradox which exploits the following weaker property of d :

$$(7') \quad \text{If } H_1, H_2 \text{ entail } E \text{ and } \Pr(H_1) > \Pr(H_2), \text{ then } d(H_1, E) > d(H_2, E).$$

- r violates even this weaker condition, but l satisfies (7').
- In this sense, Sober's resolution of Goodman's "Grue" paradox is *less* sensitive to choice of measure (*i.e.*, more robust) than Eells's is.
- And, like Eells, Sober does provide *some* reasons to prefer d to r .
- However, as I explain in my [13] and [15], these reasons (which are borrowed from Eells) are not very good reasons to prefer d to r .
- Like Earman's account of "Irrelevant Conjunction," Sober's account of "Grue" can be bolstered by providing compelling independent reasons to favor d (or l) over r (*e.g.*, see below).

Extra Slide #10: Earman on the Quantitative Problem of Old Evidence

- Earman [6, pp. 120–121] argues that quantitative Bayesian confirmation theory, together with the "radical probabilism" of Jeffrey [26] does not suffice to avoid Glymour's problem of old evidence [17, pp. 63–69].
- His argument presupposes that Bayesians use d to measure degree of confirmation, and it rests on the following fact about d :

$$(8) \quad \text{If } H \vDash E, \text{ then } \Pr(E) \approx 1 \Rightarrow d(H, E) \approx 0.$$

- This argument has two flaws. First, (8) does hold for d and r , but it does *not* hold for l or s (contrary to what Earman suggests [6, p. 243, note 8]). Second, this argument only applies to the case of *deductive evidence* ($H \vDash E$).
 - As it turns out, we can avoid Earman's objections, by using our l instead of d :
- $$(9) \quad \text{Even if } H \vDash E \text{ and } \Pr(E) \approx 1, l(H, E) \text{ can be arbitrarily large.}$$
- As Joyce [27] and Christensen [5] point out, s also satisfies (9).

References

- [1] L. Bovens, B. Fitelson, S. Hartmann, and J. Snyder, *Too odd (not) to be true? A reply to Erik J. Olsson*, *British Journal for the Phil. of Sci.* (forthcoming) <http://fitelson.org/bjps.pdf>.
- [2] L. Bovens and S. Hartmann, *Bayesian networks and the problem of unreliable instruments*, *Philosophy of Science* **69** (2002), 29–72.
- [3] D. Bradley and B. Fitelson, *Monty hall, doomsday, and confirmation*, *Analysis* (forthcoming) <http://fitelson.org/monty.pdf>.
- [4] R. Carnap, *Logical foundations of probability*, second ed., University of Chicago Press, 1962.
- [5] D. Christensen, *Measuring confirmation*, *Journal of Philosophy* **XCVI** (1999), 437–61.
- [6] J. Earman, *Bayes or bust?*, MIT Press, 1992.
- [7] E. Eells, *Rational decision and causality*, Cambridge University Press, 1982.
- [8] E. Eells and B. Fitelson, *Measuring confirmation and evidence*, *Journal of Philosophy* **XCVII** (2000), no. 12, 663–672, <http://fitelson.org/dccfin.pdf>.
- [9] ———, *Symmetries and asymmetries in evidential support*, *Philosophical Studies* **107** (2002), no. 2, 129–142, <http://fitelson.org/symmetry.pdf>.
- [10] R. Festa, *Bayesian confirmation*, *Experience, Reality, and Scientific Explanation* (M. Galavotti and

- A. Pagnini, eds.), Kluwer Academic Publishers, 1999, pp. 55–87.
- [11] B. Fitelson, *Putting the irrelevance back into the problem of irrelevant conjunction*, forthcoming in *Philosophy of Science*, <http://fitelson.org/ic.pdf>.
- [12] ———, *Review of Joyce's [27]*, *Mind* (forthcoming), <http://fitelson.org/joyce.pdf>.
- [13] ———, *The plurality of Bayesian measures of confirmation and the problem of measure sensitivity*, *Philosophy of Science* **66** (1999), S362–S378, <http://fitelson.org/psa.pdf>.
- [14] ———, *A Bayesian account of independent evidence with applications*, *Philosophy of Science* **68** (2001), no. 3, suppl., S123–S140, <http://fitelson.org/psa2.pdf>.
- [15] ———, *Studies in Bayesian confirmation theory*, Ph.D. thesis, University of Wisconsin–Madison (Philosophy), 2001, <http://fitelson.org/thesis.pdf>.
- [16] D. Gillies, *In defense of the Popper-Miller argument*, *Philosophy of Science* **53** (1986), 110–113.
- [17] C. Glymour, *Theory and evidence*, MIT Press, 1980.
- [18] I. Good, *The best explicatum for weight of evidence*, *Journal of Statistical Computation and Simulation* **19** (1984), 294–299.
- [19] ———, *Weight of evidence: a brief survey*, *Bayesian Statistics 2*, North-Holland, 1985, pp. 249–269.
- [20] ———, *A Reinstatement, in Response to Gillies, of Redhead's Argument in Support of Induction*,

- Philosophy of Science* **54** (1987), 470–72.
- [21] A. Hájek, *What conditional probabilities could not be*, *Synthese* (forthcoming), 2002.
- [22] D. Heckerman, *An axiomatic framework for belief updates*, *Uncertainty in Artificial Intelligence 2*, Elsevier, 1988, pp. 11–22.
- [23] E. Horvitz and D. Heckerman, *The inconsistent use of certainty measures in artificial intelligence research*, *Uncertainty in Artificial Intelligence 1*, Elsevier, 1986, pp. 137–151.
- [24] P. Horwich, *Probability and evidence*, Cambridge University Press, 1982.
- [25] C. Howson and P. Urbach, *Scientific reasoning: The bayesian approach*, Open Court, 1993.
- [26] R. Jeffrey, *Probability and the art of judgment*, Cambridge University Press, 1992.
- [27] J. Joyce, *The foundations of causal decision theory*, Cambridge University Press, 1999.
- [28] J. Kemeny and P. Oppenheim, *Degrees of factual support*, *Phil. of Sci.* **19** (1952), 307–324.
- [29] A. Kolmogorov, *Foundations of probability*, second english ed., AMS Chelsea Publishing, 1956.
- [30] H. Kyburg, *Subjective probability: criticisms, reflections, and problems*, *J. Philos. Logic* **7** (1978), no. 2, 157–180.
- [31] ———, *Recent work in inductive logic*, *Recent Work in Philosophy* (T. Machan and K. Lucey, eds.), Rowman & Allanheld, 1983, pp. 87–150.

- [32] I. Levi, *Changing probability judgements*, Patrick Suppes: scientific philosopher, Vol. 1, Synthese Lib., vol. 233, Kluwer Acad. Publ., 1994, pp. 87–108.
- [33] P. Maher, *Betting on theories*, Cambridge University Press, 1993.
- [34] ———, *Subjective and objective confirmation*, *Philosophy of Science* **63** (1996), 149–174.
- [35] P. Milne, $\log[p(h/eb)/p(h/b)]$ is the one true measure of confirmation, *Philosophy of Science* **63** (1996), 21–26.
- [36] H. Mortimer, *The logic of induction*, Prentice Hall, 1988.
- [37] R. Nozick, *Philosophical explanations*, Harvard University Press, Cambridge, 1981.
- [38] J. Pearl, *Probabilistic reasoning in intelligent systems: Networks of plausible inference*, Morgan Kaufman, San Francisco, 1988.
- [39] C. Peirce, *The probability of induction*, *Popular Science Monthly* **12** (1878), 705–718.
- [40] K. Popper, *The logic of scientific discovery*, Routledge, 1992.
- [41] K. Popper and D. Miller, *The impossibility of inductive probability*, *Nature* **302** (1983), 687–688.
- [42] M. Redhead, *On the impossibility of inductive probability*, *The British Journal for the Philosophy of Science* **36** (1985), 185–191.
- [43] H. Reichenbach, *The direction of time*, University of California, 1956.

- [44] P. Roeper and H. Leblanc, *Probability theory and probability logic*, University of Toronto, 1999.
- [45] R. Rosenkrantz, *Foundations and applications of inductive probability*, Ridgeview, 1981.
- [46] ———, *Bayesian confirmation: Paradise regained*, *The British Journal for the Philosophy of Science* **45** (1994), 467–476.
- [47] R.D. Rosenkrantz, *Bayesian theory appraisal: a reply to Seidenfeld*, *Theory and Decision* **11** (1979), no. 4, 441–451.
- [48] D. Schum, *The evidential foundations of probabilistic reasoning*, John Wiley & Sons, 1994.
- [49] T. Seidenfeld, *Why I am not an objective Bayesian; some reflections prompted by Rosenkrantz*, *Theory and Decision* **11** (1979), no. 4, 413–440.
- [50] G. Shafer, *A mathematical theory of evidence*, Princeton University Press, 1976.
- [51] E. Sober, *Independent evidence about a common cause*, *Philosophy of Science* **56** (1989), 275–287.
- [52] ———, *No model, no inference: A bayesian primer on the grue problem*, *Grue! The New Riddle of Induction* (D. Stalker, ed.), Open Court, 1994.
- [53] L. Wittgenstein, *Philosophical investigations*, Macmillan, 1953.
- [54] L. Zadeh, *Probability measures of fuzzy events*, *J. Math. Anal. Appl.* **23** (1968), 421–427.