

A problem for the alternative difference measure of confirmation

Nevin Climenhaga

Published online: 6 March 2012
© Springer Science+Business Media B.V. 2012

Abstract Among Bayesian confirmation theorists, several quantitative measures of the degree to which an evidential proposition E confirms a hypothesis H have been proposed. According to one popular recent measure, s , the degree to which E confirms H is a function of the equation $P(H|E) - P(H|\sim E)$. A consequence of s is that when we have two evidential propositions, $E1$ and $E2$, such that $P(H|E1) = P(H|E2)$, and $P(H|\sim E1) \neq P(H|\sim E2)$, the confirmation afforded to H by $E1$ does not equal the confirmation afforded to H by $E2$. I present several examples that demonstrate the unacceptability of this result, and conclude that we should reject s (and other measures that share this feature) as a measure of confirmation.

Keywords Confirmation · Evidence · Bayesian epistemology · Probability

In recent years there has been a debate among Bayesians as to which measure of the level of confirmation afforded to a hypothesis H by an evidential proposition E is most adequate. Following convention, I will refer to the level of confirmation afforded to H by E according to some measure as $c(H,E)$. Several measures have been defended in the literature, but in recent years four have probably been the most prominent.¹ According to the difference measure (d),

$$c(H, E) = P(H|E) - P(H)$$

¹ For example, these are the four considered in Fitelson (2001).

The alternative difference measure (s) uses the equation²

$$P(H|E) - P(H|\sim E)$$

And the ratio measure (r) and likelihood ratio (l) measures, respectively, use

$$\log\left(\frac{P(H|E)}{P(H)}\right)$$

and

$$\log\left(\frac{P(E|H)}{P(E|\sim H)}\right)$$

Though probably not as popular as the other measures,³ s has received qualified endorsement from Christensen (1999) and Joyce (1999, 2004).⁴ Both are sceptical that a single measure of confirmation will satisfy all the intuitions we have about the confirmatory relation. Christensen, however, regards s as an improvement over d , r , and l , arguing that the manner in which the alternative difference measure models confirmation (in particular, the confirmatory power of “old evidence”) is superior to the manner in which the other measures do so. Joyce (1999, p. 206), for his part, thinks that d and s both “capture contrasting, but equally legitimate, ways of thinking about evidential relevance, which serve somewhat different purposes”.⁵ The purpose of this paper is to identify a significant problem for s that has not previously been recognized—namely, that its violation of a certain plausible constraint on confirmation measures leads to markedly counterintuitive consequences in a wide range of cases.⁶ This problem, I think, makes it quite clear that s is not the one true measure of confirmation; but even if one agrees with Joyce that there are multiple legitimate measures, I will argue that the problem is serious enough to cast significant doubt on its being one of them.

As noted above, according to s , the degree of confirmation afforded to H by E is a function of the equation $P(H|E) - P(H|\sim E)$. A consequence of this is that when we have two evidential propositions, E_1 and E_2 , such that $P(H|E_1) = P(H|E_2)$, and $P(H|\sim E_1) \neq P(H|\sim E_2)$, $c(H, E_1) \neq c(H, E_2)$. Put another way, s violates the following condition:

$$c(H, E_1) > / = / < c(H, E_2) \text{ iff } P(H|E_1) > / = / < P(H|E_2).$$

Call this condition Conditional Correspondence (CC). In their criticism of s , Eells and Fitelson (2000, p. 671) note that s violates a principle equivalent to (CC);

² Or $\frac{P(H|E) - P(H)}{P(\sim E)}$. These are equivalent provided that $P(E)$ is not 1. See Christensen (1999, p. 450). I call measure s the alternative difference measure of confirmation because it was most prominently defended by Christensen (1999) as an alternative to the traditional difference measure d .

³ As evidenced, for example, by Fitelson’s exclusion of it in favor of d , r , and l when he briefly mentions ‘the three most popular Bayesian relevance measures of ... confirmation’ in Fitelson (2007, p. 478).

⁴ Earman considers, but quickly rejects, s in Earman (1992, pp. 120–121). Earman’s reason for rejecting s is addressed in both Christensen (1999) and Joyce (1999).

⁵ See also Joyce (2004, pp. 144–145), in which s is again defended as one of multiple legitimate measures.

⁶ As I will note, the problem I raise for s also affects other, less popular, measures that have been proposed, but I focus on s because it is seen as more of a live option than these other measures.

however, the only argument they give for that principle is that “many contemporary Bayesian resolutions of both the ravens paradox and the problem of evidential variety (or diversity) depend on” it. This is unfortunate, because one could conclude from the fact that s violates (CC) not that s is incorrect, but rather that (CC) is an inadequate principle, and so that these Bayesian resolutions of these problems are unsatisfactory. And, indeed, referring to the same principle in an earlier article, Fitelson (1999, p. S368) himself suggested that

Until we are given some compelling reason to prefer [measures that do not violate this principle to measures that do] we should be wary about accepting the popular quantitative resolutions of the Ravens Paradox, or the recent Bayesian accounts of the confirmational significance of evidential diversity.⁷

In more recent articles, Fitelson has given independent reasons to prefer certain of those measures, and it is not my purpose to comment on those here. Rather, I want to give an argument directly for (CC), and hence, against s , one that is, to my knowledge, novel. The following illustrations will show that in situations where E2 entails E1, violating (CC) is highly unintuitive.

Suppose that you are chasing after a murderer who has fled the scene of a crime. You suspect that he has escaped through the forest. As you are searching the forest, you happen upon the murder weapon, apparently dropped by the killer as he ran. This, of course, confirms your hypothesis that the murderer fled through the forest. But according to s , the degree of confirmation differs depending on the proposition you use to express the evidence before you. You might most naturally use a proposition like

W: The murder weapon is in the forest.

But you could also use the proposition

W₁₂: The murder weapon is in sector 12 of the forest,

where the forest is divided into one hundred equal square sectors numbered 1–100, and you are in sector 12 when you find the weapon.

Intuitively, which of these propositions you use to describe the evidence in front of you should not make a difference to the degree hypothesis H, that the murderer escaped through the forest, is confirmed. But, on measure s , it does. W₁₂ is much less probable than W; it is much less probable that the killer dropped his weapon in a small area of the forest than that he dropped it in the forest at all. W and W₁₂ are obviously both unlikely given $\sim H$; (W & $\sim H$) being true would presumably depend on the unlikely scenario of someone besides the killer putting the murder weapon in the forest. Let’s suppose also that you know that the killer is a clumsy oaf, and so that the chances are rather good that he dropped the weapon while fleeing. We can

⁷ (CC) is also invoked by Crupi et al. (2010, p. 79) and Crupi et al. (2007, p. 234). The former propose it as a basic *desideratum* on Bayesian measures of confirmation, and the latter offer the fact that their preferred measure z (a measure not mentioned above) satisfies (CC) as evidence for z . In both cases, however, the authors’ only argument for (CC) is that it is necessary for popular Bayesian analyses, such as Horwich’s (1982) solution to the ravens paradox. Steel (2003, pp. 219–220) also mentions principles equivalent to (CC), but does not argue that they are true, only that they are assumed by many Bayesians.

suppose, then, that the following probabilities obtain (before you have found the weapon):

$$P(H) = .5$$

$$P(W|H) = .7$$

$$P(W|\sim H) = .001$$

$$\begin{aligned} P(W) &= P(H)P(W|H) + P(\sim H)P(W|\sim H) = (.5)(.7) + (.5)(.001) \\ &= .35 + .0005 = .3505 \end{aligned}$$

If we suppose that the murderer (or someone else) is no more likely to have dropped the weapon in one section of the forest than another, then

$$P(W_{12}) = .3505/100 = .003505$$

$$P(W_{12}|H) = .7/100 = .007$$

$$P(W_{12}|\sim H) = .001/100 = .00001$$

So, by Bayes' Theorem:

$$P(H|W) = \frac{P(H)P(W|H)}{P(W)} = \frac{(.5)(.7)}{.3505} = \frac{.35}{.3505} \approx .9986$$

This is equal to

$$P(H|W_{12}) = \frac{P(H)P(W_{12}|H)}{P(W_{12})} = \frac{(.5)(.007)}{.003505} = \frac{.0035}{.003505} \approx .9986$$

$P(H|\sim W_{12})$, however, is over twice as large as $P(H|\sim W)$ (this, of course, is what one would expect—that the murder weapon is not in this small section of the forest matters little to your hypothesis, but learning that it's not in the forest at all ought to make you significantly less confident in your hypothesis, if you thought there was a reasonable chance the murderer would drop it while fleeing):

$$P(H|\sim W) = \frac{P(H)P(\sim W|H)}{P(\sim W)} = \frac{(.5)(.3)}{.6495} = \frac{.15}{.6495} \approx .2309$$

$$P(H|\sim W_{12}) = \frac{P(H)P(\sim W_{12}|H)}{P(\sim W_{12})} = \frac{(.5)(.993)}{.996495} = \frac{.4965}{.996495} \approx .4982$$

Using s ,

$$c(H, W) = P(H|W) - P(H|\sim W) \approx .9986 - .2309 = .7677$$

$$c(H, W_{12}) = P(H|W_{12}) - P(H|\sim W_{12}) \approx .9986 - .4982 = .5004$$

According to s , then, W supports H significantly more than W_{12} does. But again, which proposition you choose to use to express the evidence you have when you happen upon the weapon should not be relevant to the level of confirmation that evidence affords to the hypothesis that the murderer fled through the forest. Consider radioing the police chief, and telling him either that (W) 'I've found the murder weapon in the forest' or (W_{12}) 'I've found the murder weapon in sector 12

of the forest.’ According to *s*, the police chief should regard the first report as confirming *H* more than the second report. But clearly, the additional information that the murder weapon was in a specific sector of the forest does not lessen the confirmation of the hypothesis that the killer escaped through that forest.

In this example, you had an experience (finding the weapon) that it was possible for you to describe using different propositions, and because the more specific proposition had a greater probability of being false, that proposition came out, according to *s*, as less confirmatory of your hypothesis than the more general proposition. Obviously one could describe almost any experience that one takes as confirmatory of some hypothesis in different terms and get this same result. But it is not essential that *E1* and *E2* be two different ways of describing the same experience for it to be counterintuitive to find that, despite the probability of *H* conditional on *E1* being equal to the probability of *H* conditional on *E2*, *E1* and *E2* confirm *H* to different degrees. Another kind of case that brings out the counter intuitiveness of violating (CC) is one in which *E1* is relevant to *H* and *E2* is *E1* conjoined with some irrelevant independent proposition; here again, *E2* will entail *E1* but be less confirmatory than it simply by virtue of providing additional (irrelevant) information. For example, let *H* be the proposition that all of the balls in the urn in front of you are black, *B* the proposition that the ball you have just pulled out is black, and *X* the proposition that the ball you have just pulled out has a platypus on it. (*B*&*X*) entails *B*, and *X* is (presumably) independent of *H* and *B*, so that $P(H|B\&X) = P(H|B)$. Since *X* and *H* are probabilistically independent, *X* does not confirm or disconfirm *H* at all (on any measure). Intuitively, then, (*B*&*X*) confirms *H* to the same degree as *B*. But *s* must say that the confirmation by *B* is greater, since $P(H|\sim B)$ is 0, whereas (since *X* is such an unlikely scenario) $P(H|\sim (B\&X))$ is almost the same as $P(H)$.

The other standard measures mentioned above do not fall prey to this problem. Recall that *d* uses the expression $P(H|E) - P(H)$, and $r \log\left(\frac{P(H|E)}{P(H)}\right)$. Since (to return to the forest example) $P(H|W) = P(H|W_{12})$, these will come out the same whichever proposition we fill in for *E*. As for *l*, which uses the expression $\log\left(\frac{P(E|H)}{P(E|\sim H)}\right)$, these come out the same whether we’re using *W* or *W*₁₂:

$$\log\left(\frac{P(W|H)}{P(W|\sim H)}\right) = \log\left(\frac{.7}{.001}\right) = \log(700) \approx 2.8451$$

$$\log\left(\frac{P(W_{12}|H)}{P(W_{12}|\sim H)}\right) = \log\left(\frac{.007}{.00001}\right) = \log(700) \approx 2.8451$$

On any of these measures *W* and *W*₁₂ confirm *H* to the same degree, as one would expect.⁸

⁸ I should also note that Crupi et al.’s (2007) recently proposed measure *z* satisfies (CC), and so does not suffer from this problem either. According to *z*, $c(H, E) = \frac{P(H|E) - P(H)}{P(\sim H)}$ if $P(H|E) \geq P(H)$, and $\frac{P(H|E) - P(H)}{P(H)}$ otherwise.

Other less popular measures have the same problem as s , though. According to Carnap's (1962) favored measure of confirmation, $c(H,E) = P(H \& E) - P(H)P(E)$. Rewritten as $P(E)P(H|E) - P(H)P(E)$, we can see that in a situation like those described above (where $P(H|E1) = P(H|E2)$), changes in the value of $P(E)$ will lead to changes in the value of $c(H,E)$. The problem also affects Mortimer's (1988) proposed measure $P(E|H) - P(E)$ and Nozick's (1981) $P(E|H) - P(E|\sim H)$, in most situations (including those above) where $P(H|E1) = P(H|E2)$, but either $P(E1|H) \neq P(E2|H)$ or $P(E1|\sim H) \neq P(E2|\sim H)$ (or both).

The common feature of all these measures is that they satisfy a condition that Eells and Fitelson (2002) call Evidence Symmetry (ES), defined as follows:

$$c(H, E) = -c(H, \sim E)$$

Informally, this says that the degree to which E confirms H is the same as the degree to which E 's negation disconfirms H . One might defend this symmetry by thinking of E as assuaging doubts about H that arise from considering the possibility that $\sim E$.⁹ If learning E confirms H by eliminating our $\sim E$ -based doubts about H , then it makes sense that E confirms H to the same degree as $\sim E$ disconfirms it.¹⁰ Applying this to our hypothesis that the murderer escaped through the forest, if learning W makes us more confident of H because it assures us that $\sim W$ is false and learning W_{12} makes us more confident of H because it assures us that $\sim W_{12}$ is false, then it makes sense to view the former as more confirmatory, since learning $\sim W_{12}$ would not be very damaging to our hypothesis (we didn't expect the murderer to drop his weapon in sector 12, so learning that he didn't doesn't make us seriously doubt that he fled through the forest). Since $P(H)$ is more affected by $P(W)$ than $P(W_{12})$, being confident in H is more dependent on being confident in W than in W_{12} , so in this sense, one might argue, learning that W does confirm H more than learning that W_{12} .

Moreover, if Joyce is right, and there are multiple equally legitimate ways of measuring confirmation, then a defender of the alternative difference measure need not necessarily deny that W and W_{12} confirm H equally in *some* sense. An advocate of s could grant that the above criticisms of s demonstrate that it does not capture *every* sense of confirmation, but maintain that in satisfying (ES) s captures one important aspect of our thinking about confirmation, and that the verdicts it delivers in the cases discussed above are in some sense correct.

I think, though, that reflection on those cases can help us see that the above defence of (ES) is not a good one, and so cast serious doubt even on the weak claim that s captures one of many legitimate senses of confirmation. What is crucial in the forest case is that W_{12} and W are not unconnected propositions. Rather, W_{12} entails W , and you know this. Hence, you can't be confident in W_{12} without being confident in W . So W_{12} reveals the falsity of not only $\sim W_{12}$, but also $\sim W$ (i.e., in learning that the weapon is in sector 12, you learn that it is false that it is not in sector 12 *and* that it is false that it is not in the forest at all); all we need is one more

⁹ My thanks to David Christensen for suggesting this possible motivation for (ES) to me.

¹⁰ Cf. Joyce's (1999, p. 206) claim, in defence of s , that one way we think about the degree to which a proposition confirms a hypothesis is by contrasting the effects of learning that proposition with the effects of learning its negation.

logical step. W_{12} assuages not only our $\sim W_{12}$ -based doubts but *also* our $\sim W$ -based ones; hence, it ought to confirm H to the same degree that simply learning that W does. Evidence Symmetry captures the effect that our newfound confidence in E has on our confidence in $\sim E$ at the expense of all other relevant logical consequences of our confidence in E. This should make us sceptical that (ES) is satisfied by *any* legitimate confirmation measure. If learning that E is true rules out propositions besides $\sim E$, and if these other propositions may in fact be much more relevant to H than $\sim E$, then holding that $c(H,E)$ ought to equal $-c(H,\sim E)$ seems just as arbitrary as holding that $c(H,E)$ ought to equal $-c(H,(\sim E\&P))$, for any P.¹¹ ($\sim E\&P$) is falsified by E just as much as $\sim E$ is, but no one would take this latter condition seriously because $\sim E\&P$ might be comparatively irrelevant to H. As the forest scenario illustrates, however, $\sim E$ might be practically irrelevant itself, while other propositions falsified by E are quite relevant.

It seems to me, then, that Evidence Symmetry is not a feature of any intuitive sense of confirmation, and so that the quantity s measures is not one that is important from the standpoint of confirmation theory. I should also note that Eells and Fitelson (2002) have provided reasons independent of those canvassed here to doubt (ES). Combined with the unintuitive consequences of (ES) examined above, I think we have strong grounds to reject any measure of confirmation that satisfies this principle.

Nevertheless, one might still wonder if all this suffices to show that s is inadequate in *all* circumstances. After all, it was crucial to the set-up of both the forest and the urn case that the one, more specific, evidential proposition (that the weapon is in *sector 12* of the forest; that the ball is black *and has a platypus on it*) entailed the other, more general, one (that the weapon is in the forest at all; that the ball is black, whether or not it has a platypus on it). An advocate of s , then, might be tempted to suggest that even if s is not an adequate measure in the kinds of cases I have discussed, it could still be useful as a measure of relative degrees of confirmation in certain other kinds of cases. At the least, these cases will be ones where none of the evidential propositions being compared entail any of the others, but perhaps there are further conditions that need to be satisfied as well (e.g., the propositions are mutually exclusive, the propositions are independent, etc).¹²

It is admittedly difficult, once we countenance confirmational pluralism, to demonstrate that a proposed confirmation measure *never* (uniquely) measures a quantity that we would intuitively identify as (in some sense) the degree to which a particular evidential proposition confirms a hypothesis. Nevertheless, I will now

¹¹ And holding that the confirmation of H by E can be measured by the difference between $P(H|E)$ and $P(H|\sim E)$ because this satisfies (ES) seems just as arbitrary as holding that it can be measured by the difference between $P(H|E)$ and $P(H|(\sim E\&P))$. Of course, this latter function could not be a general Bayesian measure of confirmation because it might be possible that $P(H|(\sim E\&P)) > P(H|E) > P(H)$ (e.g., if P entailed H and E did not), in which case E will confirm H but the measure will report that it disconfirms it. But it is interesting to note that the “measure” $P(H|E) - P(H|\sim E)$ has the same property that Christensen and Joyce find so appealing in the measure $P(H|E) - P(H|\sim E)$, namely, it is “invariant under learning” (Joyce 1999, p. 207)—changes in $P(E)$ do not lead to changes in $c(H,E)$. This property is crucial to Christensen’s (1999, pp. 451–452) use of s to solve the old-evidence problem.

¹² My thanks to an anonymous reviewer for making me aware of the need to consider this possibility.

attempt to do just that. Let us return to our forest hypothesis, and consider the proposition.

W_{1-10} : The murder weapon is in one of sectors 1–10 of the forest.

The following relations will hold between W , W_{1-10} , and W_{12} :

$$\begin{aligned} P(H|W) &= P(H|W_{1-10}) = P(H|W_{12}) \\ P(H|\sim W) &< P(H|\sim W_{1-10}) < P(H|\sim W_{12}) \\ c(H, W) &> c(H, W_{1-10}) > c(H, W_{12}) \text{ (according to } s \text{)} \end{aligned}$$

Not only does W_{1-10} not entail W_{12} , the two propositions are mutually exclusive. W_{1-10} entails W , of course. Since this is the case, the present proposal would grant that s is not an adequate measure of the relative degrees of confirmation afforded to H by W and W_{1-10} . However, if we judge that these ought to be equal, and that $c(H, W)$ and $c(H, W_{12})$ ought to be equal, we would seem to be obliged to judge that $c(H, W_{1-10})$ and $c(H, W_{12})$ ought to be equal as well. So if it is unintuitive that W confirms H more than W_{12} , it is equally unintuitive that W_{1-10} does.

In this case W_{1-10} and W_{12} were still relevant to each other, just negatively so; but this is not necessary either for s 's violating (CC) to be counterintuitive. Consider a proposition $W_{m/n}$ that says that the murder weapon is in one of a large number of discrete locations of equal size in the forest, where n = the total number of locations, m = the number of locations that fall inside sector 12, and $m/n = P(W_{12})$ (.003505 in our original setup). Then, assuming that the weapon is no more likely to be in any one location than any other, $P(W_{12}|W_{m/n}) = P(W_{12})$, and so $W_{m/n}$ and W_{12} are probabilistically independent. However, we can let the locations specified in $W_{m/n}$ cover a larger total section of the forest than sector 12,¹³ and then, according to s , $W_{m/n}$ will confirm H more than W_{12} , despite the fact that $P(H|W_{m/n}) = P(H|W_{12})$. But obviously, for exactly the same reasons as in the other cases, $W_{m/n}$ ought not confirm H more than W_{12} . In all these cases W_{12} comes out as less confirmatory *simply in virtue of providing more information*, where this extra level of detail is not relevant to the probability of H . Whether W_{12} is being compared with an evidential proposition that it is positively relevant towards, negatively relevant towards, or independent of, the mere fact that it is more informative than that other proposition should not make it less confirmatory.

Is there then no way to restrict the application of s to avoid problematic cases of the kind I have discussed? We have seen that the direct probabilistic dependence (or lack thereof) of the two evidential propositions $E1$ and $E2$ on each other is not important to the setup of such cases. What *was* crucial in all these cases, however, was the mutual positive relevance of $E1$ and $E2$ to some third proposition itself relevant to the hypothesis H .¹⁴ (In the original cases, this third proposition was

¹³ For example, if we let each location be equal to 1/10,000 of a sector, and set n to 200,000, then the locations will cover one-fifth of the forest. (If $m/n = .003505$, m will be equal to 701, and the locations will cover 701/10,000 of sector 12.)

¹⁴ It is not necessary that this positive relevance be entailment in either case. To see this, imagine that one (but only one) of the locations in $W_{m/n}$ had been set to outside the forest. Then $W_{m/n}$ would not entail W , but it would still make W (and by extension, H) highly probable, and the counter intuitiveness of $W_{m/n}$ confirming H more than W_{12} would remain.

identical to E1.) But the defender of s cannot restrict its application only to cases where E1 and E2 are not both positively relevant to some third proposition relevant to H. If E1 and E2 confirm H at all, then they must be positively relevant to some third proposition relevant to H—namely, H itself!¹⁵ Clearly a measure of confirmation that could only be used to compare the degrees to which two propositions confirm a hypothesis when one of them does not confirm the hypothesis at all is not a useful measure of confirmation. It seems, then, that there is no viable way to restrict s 's application to avoid the problem I have discussed without giving up the measure altogether. I conclude that violating (CC) is unacceptable in general, and not just in particular sorts of cases, and that, consequently, s (and any other measure that violates (CC)) is indeed inadequate for all purposes, and ought to be rejected as a measure of confirmation.

Acknowledgments This paper had its genesis in a graduate seminar on probability at Western Michigan University in Fall 2009. I am grateful to Timothy McGrew for teaching that class and helping me think through these issues. I would also like to thank David Christensen for insightful correspondence on this project, as well as Matthew Lee and an anonymous reviewer for helpful comments on earlier drafts.

References

- Carnap, R. (1962). *Logical foundations of probability* (2nd ed.). Chicago: University of Chicago Press.
- Christensen, D. (1999). Measuring confirmation. *Journal of Philosophy*, 96, 437–461.
- Crupi, V., Tentori, K., & Gonzalez, M. (2007). On Bayesian measures of evidential support: Theoretical and empirical issues. *Philosophy of Science*, 74, 229–252.
- Crupi, V., Festa, R., & Buttasi, C. (2010). Towards a grammar of Bayesian confirmation. In M. Suárez, M. Dorato, & M. Rédei (Eds.), *Epistemology and methodology of science* (pp. 73–93). Berlin: Springer.
- Earman, J. (1992). *Bayes or bust?: A critical examination of Bayesian confirmation theory*. Cambridge: MIT Press.
- Eells, E., & Fitelson, B. (2000). Measuring confirmation and evidence. *Journal of Philosophy*, 97, 663–672.
- Eells, E., & Fitelson, B. (2002). Symmetries and asymmetries in evidential support. *Philosophical Studies*, 107, 129–142.
- Fitelson, B. (1999). The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science*, 66(supplement), S362–S378.
- Fitelson, B. (2001). A Bayesian account of independent evidence with applications. *Philosophy of Science*, 68, 123–140.
- Fitelson, B. (2007). Likelihoodism, Bayesianism, and relational confirmation. *Synthese*, 156, 473–489.
- Horwich, P. (1982). *Probability and evidence*. Cambridge: Cambridge University Press.
- Joyce, J. (1999). *The foundations of causal decision theory*. Cambridge: Cambridge University Press.
- Joyce, J. (2004). Bayesianism. In A. R. Mele & P. Rawling (Eds.), *The Oxford handbook of rationality* (pp. 132–155). New York: Oxford University Press.
- Mortimer, H. (1988). *The logic of induction*. Paramus: Prentice Hall.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge: Harvard University Press.
- Steel, D. (2003). A Bayesian way to make stopping rules matter. *Erkenntnis*, 58, 213–222.

¹⁵ Also, (E1∨E2), (H∨P) for any P with probability <1, etc.

Copyright of *Philosophical Studies* is the property of Springer Science & Business Media B.V. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.