

Bayesian Confirmation Theory

- Administrative: Please consult/collaborate with me on your papers
- Bayesian Confirmation Theory
 - Some Problems for Bayesian Confirmation Theory
 - * “Old evidence”, “new theories”, and Bayesian confirmation
 - One last comment (I can’t let this go!)
 - * The (infamous) problem of “subjectivity” (priors, etc.)
 - Where do priors come from?
 - Where did the objectivity of science go?
 - Some Bayesian responses and clarifications
 - Some “Success Stories” for Bayesianism
 - * Ravens and Grue
 - * Quine-Duhem
 - * Power of independent evidence

Old Evidence/New Theories: One Last Commentary

- The key premise in the old evidence/new theories critiques seems to be:
 - (*) A rational Bayesian agent a is *always* (i.e., at all times t) obliged to make *all* judgments in accordance with their “most recent” / “most well-informed” degrees of belief $\text{Pr}_t^a(\cdot)$.
- This assumption (*) seems to be an unfair requirement on a Bayesian statistical modeler. Shouldn’t a Bayesian be able to say things like “any evidence generated by an experiment of kind K will count as evidence in favor of H (as opposed to the alternative(s) H' to H)”?
- My initial idea (also somewhat extreme!) was to say that a rational Bayesian agent (*qua* statistical modeler) should *always* use whatever probabilities are imposed by the model \mathcal{M} they choose to use.
- There are many senses of “confirms”. In some contexts we may want “confirms” to have a more concrete, personalistic flavor, while in others, we may want “confirms” to be more “hypothetical”.

Problems of Subjectivity I

- The most infamous complaint about Bayesian confirmation is that it is “*subjective*”, and so it cannot capture “*objective*” evidential relations.
- The main worry seems to be about the “prior” probabilities of hypotheses $\text{Pr}(H)$. Where do they come from? Can agents just pick priors “at random” if they like? Are there “objective” priors?
- This is important, since judgments of degree of confirmation will depend on the values of the priors (as will the posteriors!).
- Many attempts have been made to find “objectively correct” priors, in cases where “no assumptions” are made (“informationless” priors).
- The *Principle of Indifference* (or *Insufficient Reason*) says that if one has “no relevant information”, then one should assign *equal* probabilities to all (mutually exclusive, exhaustive) possibilities.
- PI is *very difficult* to apply consistently. See VF’s *Laws and Symmetry*.

Problems of Subjectivity II

- No satisfactory account of “informationless (or objective) priors” has been formulated (see B & S’s *Bayesian Theory* and VF’s *Laws and Symmetry*).
- Most Bayesians have given up on the “holy grail” of objective or informationless priors. *Merger of opinion* results are a much more popular way of restoring “objectivity” (or “intersubjectivity”).
- It can be shown (see Earman’s chapter 6) that (given weak assumptions about the process) the opinions of a community of Bayesian agents will “merge in the long run,” as more and more data are collected.
- Even if Bayesians begin with radically different priors, these priors will eventually be “washed out” by incoming evidence (e.g., 10^{100} tosses).
- There are problems with these results (see pages 147–149 of Earman).
 - There is no way to tell (generally) how long the “long run” is.
 - What about degrees of support *now*?

“Success Stories” of Bayesian Confirmation I

- It is sometimes claimed (see Earman page 64, sort of tongue-in-cheek) that Bayesian confirmation is able to “winnow a valid kernel of” previous (deductive) accounts of confirmation “from their chaff.”
- Earman’s cavalcade of “success stories” of Bayesian confirmation is fascinating. This list includes:
 - How Bayesian confirmation nicely generalizes H-D
 - How BC handles the paradoxes of instance confirmation (ravens, etc.)
 - How BC handles “evidential variety” or “diversity”
 - How BC handles the Quine-Duhem problem
 - How BC handles “grue” like paradoxes
- I will discuss some of these (leaving many details for possible paper topics and exercises — see Earman for many good hints!).

“Success Stories” of Bayesian Confirmation II

- The most famous problem in confirmation is Hempel’s paradox of the ravens. Recall this paradox was that both $E = Ra \ \& \ Ba$ and $E' = \sim Ba \ \& \ \sim Ra$ confirm $H = (\forall x)(Rx \rightarrow Bx)$.
- One key difference between Bayesian and deductive accounts of confirmation is that Bayesian accounts are *contextual* (or *three-place*). That is, we must always say “ E confirms H , relative to K .”
- IJ Good argues that, relative to some *conceivable* (but non-actual!) K ’s, $Ra \ \& \ Ba$ may not even confirm that all ravens are black!
 - $K :=$ “We are in one of two worlds: (w_1) with 100 black ravens, no nonblack ravens, and 1 million other birds, or else (w_2) with 1,000 black ravens, 1 white raven, and 1 million other birds.”
 - A bird a is selected at random from all the birds, and it is seen that $Ra \ \& \ Ba$. This seems to *disconfirm* that all ravens are black!

“Success Stories” of Bayesian Confirmation III

- Let’s assume we’re not in one of Good’s pathological worlds, and that black ravens do confirm that all ravens are black.
- The Bayesian strategy is to show that — even if a nonblack nonraven confirms that all ravens are black — it confirms that all ravens are black *much less strongly* than a black raven does.
- $c(H, E \mid K) \gg c(H, E' \mid K)$, where K is our *actual* background knowledge.
- There are *many* different Bayesian accounts of this form. Earman surveys several of these, and presents his own “clarified” version.
- I will briefly sketch an account due to Ellery Eells (in the course reader, taken from his book *Rational Decision and Causality*).
- Eells makes three assumptions about K (commonly made in here)

$\Pr(Ra \mid H \ \& \ K) = \Pr(Ra \mid K)$	$\Pr(\sim Ra \mid \sim Ba \ \& \ K) \approx 1$	$\Pr(Ba \mid Ra \ \& \ K) \not\approx 1$
--	--	--

“Success Stories” of Bayesian Confirmation IV

- Eells then adopts the *difference measure* of degree of confirmation. That is, he uses $d(H, E \mid K) = \Pr(H \mid E \ \& \ K) - \Pr(H \mid K)$.
- Eells then argues that, given his three assumptions, we have:

$$d(H, E \mid K) = \Pr(H) \cdot \left(\frac{1}{\Pr(Ba \mid Ra \ \& \ K)} - 1 \right) \\ > \Pr(H) \cdot \left(\frac{1}{\Pr(\sim Ra \mid \sim Ba \ \& \ K)} - 1 \right) = d(H, E' \mid K)$$

- This is exactly the result we wanted: that a black raven confirms that all ravens are black *more strongly* than a nonblack nonraven does.
- I think Eells’ account is one of the simpler accounts out there. But, it seems to me that it is not entirely rigorous. There seem to be gaps in parts of his derivation (Ellery intended this — it was only intended as a “sketch”). **Paper Topic:** make Eells’ argument “sketch” fully rigorous.

“Success Stories” of Bayesian Confirmation V

- It is often said that “novel” or “surprising” evidence should provide better support than “unsurprising” or “expected” evidence.
- Bayesianism gives us a way to explain this intuition. On a Bayesian account, if evidence is *improbable*, then it is also *surprising* (if learned).
- So, it would be nice if we could show that (other things being equal) *improbable evidence confirms a hypothesis more strongly than probable evidence does*. Or, more formally and generally, we’d like to show that:

(*) If $\Pr(E | K) < \Pr(E' | K)$, then $c(H, E | K) > c(H, E' | K)$.
- Earman (page 64) explains that (*) *does* hold, assuming (i) that E is *deductive evidence* (i.e., that $H \& K \models E$), and (ii) that we use the *difference measure* d as our measure c of degree of confirmation.
- Does (*) hold *more generally*, and/or for *other* measures of support c ?

“Success Stories” of Bayesian Confirmation VI

- Quine and Duhem taught us that no theory H makes predictions without substantive auxiliary assumptions K . As we saw, this caused a serious problem for the H-D account of *disconfirmation*.
- There are cases in $H \& K$ is disconfirmed by E , but we think (intuitively) that K is *more strongly disconfirmed* than H . Deductive accounts of confirmation cannot help us with this intuition.
- Dorling and, more recently (and better!) our own Michael Strevens have shown that Bayesianism can explain and model such intuitions.
- Dorling gives an historical example in which the degrees of confirmation (allegedly!) fit this pattern. Dorling makes some strong assumptions, and he reports precise numbers for the degrees of belief of historical scientists (see Earman and Strevens for criticisms).
- Strevens provides a general Bayesian *framework* for modeling Q-D.


Independent Evidence I : Wittgenstein’s Example

- Wittgenstein (in *PI*) 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 .

The Probabilistic/Causal Structure of Wittgenstein’s Example

- 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 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 explains, the relevant probabilistic fact is that E and E' are independent — *given* H . That is, if we know the truth-value of H , then the dependence (correlation) between E and E' disappears.
- In other words, H *explains* the correlation between E and E' . This is because E and E' are *joint effects of the common cause* H . In Salmon’s terminology, H, E , and E' form a *conjunctive fork*.


Picturing the Structure of the Two Examples




- 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 \bar{H}).

Two Principles About Independent Evidence

- 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 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:



Measures of Support & Independent Evidence

- The only measure (up to ordinal equivalence) that (generally) satisfies both of these principles about independent evidence is l .
- This provides a novel way of adjudicating between d and l (although d , l and r all satisfy the first, weaker principle).
- These ideas about independent evidence can also serve to ground a novel 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 .
- According to l , strong independent confirmational boosts can be provided — even to hypotheses which are highly probable *a priori*.