Bayesian Confirmation Theory

- Administrative: Almost everyone has either seen me or scheduled an appointment with me to discuss paper topics (please do) ...
- Bayesian Confirmation Theory
  - Basic Assumptions of Bayesian Epistemology
    - Degrees of belief of rational agents are *probabilities*
    - Rational degrees of belief are updated *via conditionalization*
    - 'Coherence' (probabilism) is *only* constraint on rationality
  - Two Kinds of (Qualitative) Bayesian Confirmation
  - Quantitative Bayesian Confirmation — Measures of Support
  - Some Problems for Bayesian Confirmation Theory
    - "Old evidence", "new theories", and Bayesian confirmation
  - Time Permitting: Some Success Stories for Bayesianism

Bayesian Epistemology I

- An epistemically rational agent's degrees of belief (or degrees of credence) are *probabilities* (in the technical sense).
- An (epistemically) rational agent learns by *conditionalizing* on their (total) evidence (up to that point in time).
- That is, at any given time t, a rational Bayesian agent a has degrees of belief that conform to a probability function $\Pr_t^a(\cdot)$.
- And, if a rational Bayesian agent a learns E between t_1 and t_2, then their degrees of belief at t_2 are given by $\Pr_{t_2}^a(\cdot) = \Pr_{t_1}^a(\cdot \mid E)$.
- Moreover, it is assumed that a Bayesian agent (scientist) has a *complete* probability distribution over *all* competing hypotheses.
- That is, a Bayesian agent will have precise probabilities – both *prior* probabilities $\Pr(H)$ and *posterior* probabilities $\Pr(H \mid E)$.

Bayesian Epistemology II

- The "priors" are really conditionalized on an agent's *background knowledge* K (here, K plays the role of Ω in the formal theory).
- The background knowledge (or, contextual assumptions) K plays a crucial role. Altering K can change *all* probabilistic judgments.
- Bayesian epistemology does *not* require that degrees of belief "correspond" to the "objective" probabilities (if there be such!).
- All that is required is *coherence* (that is, conformity to the probability axioms, and to the rule of conditionalization).
- Here is a classic quote from De Finetti on this issue:

  > By denying any objective value to probability I mean to say that, however an individual evaluates the probability of a particular event, no experience can prove him right, or wrong; nor, in general, could any conceivable criterion give any objective sense to the distinction one would like to draw, here, between right and wrong.

Bayesian Confirmation I

- Historically, there have been two kinds of Bayesian confirmation:
  1. **Absolute**: E confirms H (relative to K) if $\Pr(H \mid E \,\&\, K) > \tau$, for some "threshold value" τ (*i.e.*, if $\dfrac{E \,\&\, K}{\therefore H}$ is "*inductively strong*").
  2. **Incremental**: E confirms H (rel. to K) if $\Pr(H \mid E \,\&\, K) > \Pr(H \mid K)$ (*i.e.*, if E is *positively stochastically relevant to* H, given K).

  Incremental confirmation has become more popular in recent years. It will be the main focus of our discussion of Bayesian confirmation.

- Absolute confirmation is just Skyrms' account of "inductive strength".
- Incremental confirmation is probabilistic *relevance*. In this sense, it is more like the notion of "corroborative evidence" in Skyrms' chapter 8.
- These accounts differ dramatically (as we've seen in our discussion of Skyrms). Both accounts also differ greatly from *deductive* accounts.

Bayesian Confirmation II

- Some differences between the two Bayesian accounts:
  - Absolute confirmation is *insensitive to relevance* (Fred Fox, etc.).
  - Absolute confirmation is *a*symmetric, incremental is *symmetric* (here, symmetry means X confirms $Y \Longrightarrow Y$ confirms X).
  - Absolute confirmation is the *only* account we've seen with the following property: E confirms $H \not\Rightarrow \sim E$ does *not* confirm H.
  - Absolute account satisfies (SCC), incremental account does not.
- Some differences between Bayesian and deductive account:
  - Incremental account satisfies *neither* (SCC) *nor* (CCC).
  - *Neither* Bayesian account satisfies the *consistency condition* (CC).
  - *Both* bayesian accounts easily admit of *degrees*.
  - *Both* Bayesian accounts make sense in *statistical* settings.

Bayesian Confirmation III

- The incremental account will be our main focus. [It seems to me that any account which eschews *relevance* cannot be adequate.]
- So far, the incremental theory just gives us a *qualitative* account of confirmation (analogous to the previous, *deductive* accounts).
  - E *confirms* H (relative to K) if $\Pr(H \mid E \,\&\, K) > \Pr(H \mid K)$
  - E *dis*confirms H (relative to K) if $\Pr(H \mid E \,\&\, K) < \Pr(H \mid K)$
  - E is *irrelevant to* H (rel. to K) if $\Pr(H \mid E \,\&\, K) = \Pr(H \mid K)$
- The basic idea behind the incremental account is that *confirmation is probabilistic relevance*. How can we *quantitatively* generalize this?
- Each of these says "E incrementally confirms H relative to K":
  - $\Pr(H \mid E \,\&\, K) > \Pr(H \mid K)$
  - $\Pr(H \,\&\, E \mid K) > \Pr(H \mid K) \cdot \Pr(E \mid K)$
  - $\Pr(E \mid H \,\&\, K) > \Pr(E \mid \bar{H} \,\&\, K)$

Bayesian Confirmation IV

- *Prima Facie*, any of these inequalities (or any other equivalent inequality!) could be used to generate a (initially plausible) *quantitative measure* of *degree* of (incremental) confirmation.
- For instance, we might adopt one of the following three measures:

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

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

$$\begin{aligned} l(H, E \mid K) &=_{df} \log\left[\frac{\Pr(E \mid H \,\&\, K)}{\Pr(E \mid \bar{H} \,\&\, K)}\right] \\ &= \log\left[\frac{\Pr(H \mid E \,\&\, K) \cdot [1 - \Pr(H \mid K)]}{[1 - \Pr(H \mid E \,\&\, K)] \cdot \Pr(H \mid K)}\right]. \end{aligned}$$

- The only reason we take *logarithms* in the definitions of r and l is to ensure that (like d) r and l are (i) *positive* for *confirmation*, (ii) *zero* for *irrelevance*, and (iii) *negative* for *dis*confirmation.

Bayesian Confirmation V

- These and many other measures have been proposed and defended in the literature on (incremental) Bayesian confirmation theory.
- In my dissertation (and in several publications), I survey the many measures that have been proposed and show that a wide variety of arguments in the literature depend on which measure one chooses.
- It turns out (somewhat surprisingly) that one's choice of incremental measure of confirmation has *great* impact on one's analyses.
- This is because the truth-value of "E_1 confirms H_1 more strongly than E_2 confirms H_2" varies greatly (and in crucial ways!) depending on which measure of incremental confirmation one uses.
- How can this be, if the measures capture the same *qualitative* notion?
- Some of the new paper topics involve exploiting this *sensitivity to choice of measure of confirmation* (*e.g.*, Earman on old evidence).

Bayesian Confirmation VI

- The absolute account of confirmation has (potential) problems:
  - The problem of the threshold value τ (what should τ be?)
  - The problem of insensitivity to *relevance* (Fred Fox, Skyrms)
- The incremental account of confirmation has (potential) problems:
  - The problem of measure sensitivity (which I have made famous!)
  - The problem of old evidence (also a problem for absolute — *why?*)
- The problem of old evidence has caused a great stir in the literature. One of the new paper topics concerns this famous problem.
- If a Bayesian agent is always supposed to make judgments according to their "most recent" or "most well-informed" probability function Pr, then Pr will *include* all of their evidence E up to that point. But, then $\Pr(H \mid E) = \Pr(H)$, and so *their "old" evidence is no longer evidence*!

Bayesian Confirmation VII

- There are *many* proposed resolutions of the old-evidence problem:
  - **Historical**. If the question is "Does E confirm H relative to our *current* (actual) K?", then we need to look at $\Pr(H \mid E \,\&\, K')$ *vs* $\Pr(H \mid K')$, where K' is "the K we had just prior to learning E.
  - **Counterfactual**. Similar to historical, but K' is "the K we *would have* had, had we *not* learned E" and $\Pr(H \mid E \,\&\, K')$ is "the probability we would have assigned H had we *then* learned E".
  - **Logical Relation**. Rational Bayesian agents need *not* be logically omniscient, and so an agent may learn $H \vdash E$. By assuming various things about how probabilities get assigned to "$H \vdash E$", Garber *et al* show that agents can have $\Pr(H \mid H \vdash E) > \Pr(H)$, even if E is "old evidence". Earman (chapter 5) has a great discussion on this.
- The logical relation approach (unlike the others) can also cope with the problem of *new theories* (H *not even pondered* before E learned).

Bayesian Confirmation VIII

- Like everybody else, I have an opinion (in formation!) about these thorny problems of old evidence and new theories.
- My idea (today!): Let's look at how Bayesian statisticians *actually make* inferences in cases of "old evidence" and/or "new theories".
- In their encyclopedic text *Bayesian Theory*, the prominent Bayesian Statisticians Bernardo & Smith give us a hint (their emphasis):

  > One further point about the terms prior and posterior is worth emphasizing. *They are not necessarily to be interpreted in a chronological sense*, with the assumption that 'prior' beliefs are specified first and then later modified into 'posterior' beliefs. ... It is true that the natural order of assessment does coincide with the 'chronological' order in a number of applications, but ... this is a pragmatic issue and not a requirement of the theory.

- How are statistical inferences actually made? What do they establish? I think careful attention to these questions may yield a new (and superior) resolution of the old evidence and new theory problems.

Bayesian Confirmation IX

- Aren't statistical inferences really of the form "E favors H_1 over $\{H_2, \ldots, H_n\}$, *relative to an experimental design (and context) K*"?
- If so, then why can't such claims be true in a *timeless* way? And, why should they be *undermined* if E is "currently known" or if a new hypothesis H^* is "currently out there, but as yet undiscovered"?
- This idea is *neither* "historical" *nor* "counterfactual." *Why?*
- **Analogy**: When we say that an argument $\frac{p}{\therefore q}$ is *inductively strong*, we must be clear that we mean *inductively strong* — *relative to background* K. If K contains information that *undermines* or *defeats* the inference from p to q, then — relative to K — this is *not* a strong inference.
- What this means is that inductive strength is *indexical* (or *contextual*). But, why (as B & S might ask) must that render it *chronological*?
- Can you apply this idea to the "old evidence"/"new theory" problem?

Bayesian Confirmation X

- 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 (raven, etc.)
  - How BC handles "evidential variety" or "diversity"
  - How BC handles the Quine-Duhem problem
  - How BC handles "grue" like paradoxes
- But, BC also faces its own peculiar challenges, like the problems of subjectivity (especially, for priors), old evidence/new theories, Popper-Miller, measure-sensitivity, zero priors, and various others.