

Philosophy 148 — Announcements & Such

- HW #4 grades posted ($\mu = 75$). [This one was tougher than I thought.]
- New Plan for HW #5 (owing to my flu)
 - It will be due on the last day of class — next Thursday 5/8.
 - Our discussion for it will be next Tuesday 5/6 @ 6pm (room TBA).
- I will also be preparing some final extra-credit problems. They will be distributed next week, and due at the final exam (5/20 @ 8am).
- The final exam is **Tuesday, May 20 @ 8am @ 20 Barrows**.
 - I will hold a review session the day before the final (May 19). Would a time in the afternoon (say 4-6pm) work for people?
- Today's Agenda
 - The Raven Paradox (cont'd)
 - Next: The Grue Paradox

Digression on Bayesian Confirmation Theory 8

- As we have already seen for Carnap's confirms_i, Bayesian confirmation theory will accept some of Hempel's desiderata, and reject others.
 - The EQC, the EC, and the NTC all seem quite intuitive, and they are satisfied by any probabilistic account of confirmation.
 - CC is *not* intuitive. Typically, competing theories are *not* consistent (they're mutually exclusive). Let K describe the typical probability model of a standard deck of 52 cards. Then, consider the following examples.
 - E = card is black, H = card is the A♠, H' = card is the J♣. Intuitively, E confirms both H and H' (rel. to K), even though they are inconsistent.
 - SCC is not intuitive either. Many intuitive counterexamples are out there. *E.g.*, let E = card is black, H = card is the A♠, and H' = card is an ace.
-
- As for CCC, it is *highly unintuitive* (here, we agree with Hempel). *E.g.*, let E = card is the A♠, H = card is card is an ace, and H' = card is the A♦.

Qualitative, Comparative, and Quantitative Confirmation

- E confirms H , relative to K iff
 - Nicod: E is a *syntactical* instance of H (only applies to universal H s).
 - Hempel: $E \& K \models \text{dev}_{I(E)}(H)$ [$I(E)$ = individuals mentioned in E].
 - H-D: $H \& K \models E$.
 - Absolute Bayes (confirms_f): $\Pr(H | E \& K) > r$, for some “threshold” r .
 - Incremental Bayes (confirms_i): $\Pr(H | E \& K) > \Pr(H | K)$.
- E confirms H rel. to K more strongly than E' confirms H' rel. to K' iff
 - Nicod, Hempel, H-D: No accounts (purely qualitative theories).
 - Absolute Bayes: $\Pr(H | E \& K) > \Pr(H' | E' \& K')$.
 - Incremental Bayes: $c(H, E | K) > c(H', E' | K')$ [some \mathfrak{X} -measure c].
- E confirms H relative to K to degree α :
 - Nicod, Hempel, H-D: No accounts (purely qualitative theories).
 - Absolute Bayes: $f(\Pr(H | E \& K), r) = \alpha$ [some function f , thresh. r]
 - Incremental Bayes: $c(H, E | K) = \alpha$ [some \mathfrak{X} -measure c].

... back from our digression to Bayesian responses to the raven paradox ...

- Bayesians have said a great many (wildly different) things about Hempel's paradox. Almost all Bayesian approaches fall into at least one of the following three categories [H is short for “ $(\forall x)(Rx \supset Bx)$ ”, “ E_1 ” for “ $Ra \& Ba$ ” (some a drawn from the universe), and “ E_2 ” for “ $\sim Ba \& \sim Ra$ ”]:
 - **Qualitative.** Reject some precise, Bayesian rendition of (NC), on Bayesian grounds. Strictly speaking, this does not *require* the rejection of the corresponding rendition of (PC) [but this is often rejected too].
 - **Comparative.** Argue that $c(H, E_1 | K_\alpha) > c(H, E_2 | K_\alpha)$, for our *actual* background knowledge K_α . Traditionally, these approaches *accept* (PC) and do *not* deny (NC) — they *entail* $c(H, E_1 | K_\alpha) > c(H, E_2 | K_\alpha) > 0$.
 - **Quantitative.** Argue that $c(H, E_2 | K_\alpha)$ is “minute”, for our *actual* background knowledge K_α . Traditionally, these approaches go hand in hand with the comparative approaches. They typically aim to show *both* that $c(H, E_1 | K_\alpha) \gg c(H, E_2 | K_\alpha) > 0$ *and* that $c(H, E_2 | K_\alpha) \approx 0$.
- Next, I'll discuss the tradition, and then describe a new approach.

- All Bayesian approaches begin by *precisifying* (NC) [and (PC)].
- Since Bayesian confirmation is a *three-place* relation [$\mathcal{C}(H, E | K)$], we'll need a *quantifier* over the *implicit K's* in (NC). Four renditions:

$$(NC_w) \quad (\exists K)(\forall F)(\forall G)(\forall x)[\mathcal{C}((\forall x)(Fx \supset Gx), Fx \& Gx | K)]$$

$$(NC_\alpha) \quad (\forall F)(\forall G)(\forall x)[\mathcal{C}((\forall x)(Fx \supset Gx), Fx \& Gx | K_\alpha)]$$

$$(NC_T) \quad (\forall F)(\forall G)(\forall x)[\mathcal{C}((\forall x)(Fx \supset Gx), Fx \& Gx | K_T)]$$

$$(NC_s) \quad (\forall K)(\forall F)(\forall G)(\forall x)[\mathcal{C}((\forall x)(Fx \supset Gx), Fx \& Gx | K)]$$

- (NC_w) is *too weak* [let $K =$ "all instances confirm all generalizations"].
- Hempel's "explaining away" *suggests* (NC_s) should be *too strong*.
- So (NC_α) and (NC_T) seem to be the salient renditions of (NC).
- *Qualitative* Bayesians seek to refute *some* rendition of (NC).
- The question for qualitative approaches is *which* (NC) to refute.
- Early qualitative Bayesian approaches took aim at (NC_s).

- Here's Good's attempt to meet Hempel's Challenge about (NC_T):

... imagine an infinitely intelligent newborn baby having built-in neural circuits enabling him to deal with formal logic, English syntax, and subjective probability. He might now argue, after defining a crow in detail, that it is initially extremely unlikely that there are any crows, and therefore that it is extremely likely that all crows are black. ... On the other hand, if there are crows, then there is a reasonable chance that they are a variety of colours. Therefore, if I were to discover that even a black crow exists I would consider $[H]$ to be less probable than it was initially.

- Even Good wasn't so confident about this "counterexample" to (NC_T). Maher argues this is *not* a counterexample to (NC_T).
- However, Maher has recently provided a very compelling (Carnapian) counterexample to (NC_T), which is beyond our scope (unicorns).
- Most modern Bayesians don't *understand* (NC_T). Unlike Carnap, they have *no theory* of "Pr_T." They've nothing to say about (NC_T). This is why they *abandon qualitative* approaches in favor of the *comparative/quantitative*.

- I.J. Good showed that the strong (Bayesian) rendition (NC_s) of Nicod's condition is false. He gave the following counterexample:

K : Exactly one of the following two hypotheses is true: (H) there are 100 black ravens, no nonblack ravens, and 1M other birds, or ($\sim H$) there are 1,000 black ravens, 1 white raven, and 1M other birds.

E : $Ra \& Ba$ (a is randomly sampled from the universe).

So, $H = (\forall x)(Rx \supset Bx)$, and $E = Ra \& Ba$. And, we have:

$$\Pr(E | H \& K) = \frac{100}{1000100} \ll \frac{1000}{1001001} = \Pr(E | \sim H \& K)$$

- So, (NC_s) is false, and *even for "natural kinds"* (pace Quine). Similar examples can be generated to show that (PC_s) and Scheffler's (*_s) are false.
- So? Hempel replies that (NC_T) *not* (NC_s) is the salient rendition. That's plausible, but as we have seen *it's incompatible with Hempel's theory!*
- Nonetheless, Good later tried to meet Hempel's (NC_T) challenge.
- He gave the following example, which is known as "Good's Baby"...

- There have been *many* comparative Bayesian approaches. Here is the canonical, contemporary, comparative approach.
- Assume that our *actual* background corpus K_α is such that:

$$(1) \Pr(\sim Ba | K_\alpha) > \Pr(Ra | K_\alpha)$$

$$(2) \Pr(Ra | H \& K_\alpha) = \Pr(Ra | K_\alpha) \quad [\because \sim Ra \perp H | K_\alpha (!)]$$

$$(3) \Pr(\sim Ba | H \& K_\alpha) = \Pr(\sim Ba | K_\alpha) \quad [\because Ba \perp H | K_\alpha (!)]$$

Theorem. Any Pr satisfying (1), (2) and (3) will also be such that:

$$(4) \Pr(H | Ra \& Ba \& K_\alpha) > \Pr(H | \sim Ba \& \sim Ra \& K_\alpha).$$

- By (*), the observation of a black raven \therefore confirms that all ravens are black *more strongly than* the observation of a white shoe, *provided*:
 - (1) there are (proportionally) fewer ravens than non-black things
 - (2)/(3) whether something a (sampled at random from the universe) is a raven/black is *independent* of whether all ravens are black $[H]$
- As it turns out, assumptions (1)–(3) entail *a lot more than* just (4) ...

- Assumptions (1)–(3) *also* entail:
 - (5) $\Pr(H \mid Ra \ \& \ Ba \ \& \ K_\alpha) > \Pr(H \mid K_\alpha)$
 - (6) $\Pr(H \mid \sim Ba \ \& \ \sim Ra \ \& \ K_\alpha) > \Pr(H \mid K_\alpha)$
 - (7) $\Pr(H \mid Ba \ \& \ \sim Ra \ \& \ K_\alpha) < \Pr(H \mid K_\alpha)$
- I.e.*, the canonical *comparative* Bayesian assumptions entail the *qualitative* claims that $Ra \ \& \ Ba$ and (PC) $\sim Ba \ \& \ \sim Ra$ confirm H , which is *inconsistent* with denying the salient instances of (NC_α) . And, (7) is pretty crazy, no?
- So (i) this cannot undergird a *purely comparative* approach – one that is *consistent with* a qualitative $_\alpha$ approach, and (ii) it entails (7), which is bad.
- A canonical *quantitative* approach is had by *strengthening* (1) to:
 - 1'. $\Pr(\sim Ba \mid K_\alpha) \gg \Pr(Ra \mid K_\alpha)$
 - (1')–(3) $\Rightarrow \Pr(H \mid Ra \ \& \ Ba \ \& \ K_\alpha) \gg \Pr(H \mid \sim Ba \ \& \ \sim Ra \ \& \ K_\alpha) > \Pr(H \mid K_\alpha)$, and $\Pr(H \mid \sim Ba \ \& \ \sim Ra \ \& \ K_\alpha) \approx \Pr(H \mid K_\alpha)$, *i.e.*, $c(H, \sim Ba \ \& \ \sim Ra \mid K_\alpha) \approx 0$.
- That's the traditional comparative/quantitative Bayesian story.
- It would be nice to have a *purely comparative* approach . . .

- Most people seem to think that (1) – and even (1') – is true. That is, most people accept that there are (proportionally) fewer (even far fewer) ravens than non-black things in the universe (and this seems right to me).
- The controversial assumptions are the *independences*: (2) and (3).
- Vranas (on website) provides compelling objections to (2) and (3), and their standard rationales. He also argues that (3) is “for all practical purposes *necessary*” for the traditional *quantitative* conclusions.
- Pace* Vranas, it turns out that assumptions *much weaker than* (1)–(3) will suffice *both* for comparative *and* for quantitative approaches.
- (2) and (3) can be replaced by the following, weaker assumption:
 - (\ddagger) $\Pr(H \mid Ra \ \& \ K_\alpha) \geq \Pr(H \mid \sim Ba \ \& \ K_\alpha)$ [*i.e.*, $\frac{\Pr(Ra \mid H \ \& \ K)}{\Pr(Ra \mid K)} \geq \frac{\Pr(\sim Ba \mid H \ \& \ K)}{\Pr(\sim Ba \mid K)}$]
 - \Rightarrow (1) and (\ddagger) $\Rightarrow \Pr(H \mid Ra \ \& \ Ba \ \& \ K_\alpha) > \Pr(H \mid \sim Ba \ \& \ \sim Ra \ \& \ K_\alpha)$.
- (1) and (\ddagger) are consistent with denying each of the three qualitative claims (5)–(7). Thus, we have a purely comparative approach, which is (i) *consistent with* qualitative $_\alpha$ approaches, and (ii) avoids (7) as well.

- Most contemporary Bayesians are “Hempelian” in that they *accept* (PC).
- Hempel appeals to “tautological” vs “nontautological” C in his “explanation away.” But, this *contradicts* his (*monotonic*) theory of \mathcal{C} .
- This distinction *can* be theoretically grounded by *Bayesians*, *via*:
 - (\star) $c(H, \sim Ba \ \& \ \sim Ra \mid T) > c(H, \sim Ba \mid \sim Ra) = 0$
- Maher shows (\star) is true relative to a class of Carnapian \Pr_T models.
- Maher is in the minority. Most Bayesians *try* to distinguish qualitative vs comparative/quantitative. But, their strong assumptions *conflate* them.
- \Rightarrow Assumptions (1)–(3) entail that (\star) is *false!* As a result, the traditional Bayesian approach can't recover Hempel's central intuition! Bad.
- Our (1) + (\ddagger) approach $\not\Rightarrow$ (\star) is false! Thus, our approach:
 - is *consistent with* qualitative $_\alpha$ approaches.
 - avoids independencies and other bad consequences like (7).
 - is consistent with the truth of Hempel's (\star).

- Overview: Does $\sim Ba \ \& \ \sim Ra [E]$ confirm $(\forall x)(\sim Bx \supset \sim Rx) [H]$?
 - **Hempel**: Yes, relative to T. But, don't conflate this with the claim (PC*) of confirmation relative to $\sim Ra$, which is *intuitively* false. That's a nice intuition, but, unfortunately, it contradicts Hempel's theory [(M_K)].
 - **Scheffler**: Yes, but this does not imply that E confirms $(\forall x)(Rx \supset Bx)$, since E is also a positive instance of the *contrary* of $(\forall x)(Rx \supset Bx)$.
 - **Quine**: No [relative to T] because (NC_T) *does not apply* to $\sim Ba, \sim Ra$, since they are not “natural kinds,” despite the fact that Ra and Ba are “NKs”. For Quine, “NKs” must have “sufficiently similar instances”.
 - **Maher** (2004): No, not even relative to T, since (NC_T) is demonstrably *false* within a Carnapian theory of “confirmation relative to T”. Note: the falsity of (NC_T) does not depend on “naturalness” of F and G .
 - **Bayesians**: Depends on whether $\Pr(H \mid E \ \& \ K_\alpha) > \Pr(H \mid K_\alpha)$, where K_α is our *actual* background knowledge. Bayesians think that (NC_T) is *irrelevant*, epistemically, and so *they don't care* whether it's true. And, even if the K_α version is true, we can still give a *comparative* account.

Goodman's "Grue" Paradox: Basic Linguistic Structures and Facts I

- Let $Ox \stackrel{\text{def}}{=} x$ is observed prior to t , $Gx \stackrel{\text{def}}{=} x$ is green, and $Bx \stackrel{\text{def}}{=} x$ is blue.
- "Grue": $Gx \stackrel{\text{def}}{=} x$ is either observed prior to t and green or x is not observed prior to t and blue. That is, $Gx \stackrel{\text{def}}{=} (Ox \ \& \ Gx) \vee (\sim Ox \ \& \ Bx)$.
- We can also define "Bleen" as: $Bx \stackrel{\text{def}}{=} (Ox \ \& \ Bx) \vee (\sim Ox \ \& \ Gx)$.
- **Two Facts.**
 - Gx is logically equivalent to $(Ox \ \& \ Gx) \vee (\sim Ox \ \& \ Bx)$.
 - Bx is logically equivalent to $(Ox \ \& \ Bx) \vee (\sim Ox \ \& \ Gx)$.
- So, from the point of view of the Green/Blue language, "Grue" and "Bleen" are "gerrymandered" or "positional" or "non-qualitative".
- But, from the point of view of the Grue/Bleen language, "Green" and "Blue" are "gerrymandered" or "positional" or "non-qualitative".
- So, no appeal to syntax will forge an asymmetry here, unless one assumes a *privileged language*. Note: the languages are expressively equivalent.

Goodman's "Grue" Paradox: Basic Linguistic Structures and Facts II

- I'm going to simplify things by re-defining "grue" using green and non-green. Quine wouldn't have liked this, but Goodman/Hempel wouldn't have minded. It will make the subsequent discussion easier.
- Thus, "Grue" becomes: $Gx \stackrel{\text{def}}{=} Ox \equiv Gx$. Now, consider the following two universal generalizations, and three singular evidential claims:
 - H_1 : All emeralds are green. $(\forall x)(Ex \supset Gx)$.
 - H_2 : All emeralds are grue. $(\forall x)(Ex \supset Gx)$. *I.e.*, $(\forall x)[Ex \supset (Ox \equiv Gx)]$.
 - E_1 : a is a green emerald. $Ea \ \& \ Ga$.
 - E_2 : a is a grue emerald. $Ea \ \& \ Ga$. *I.e.*, $Ea \ \& \ (Oa \equiv Ga)$.
 - \mathcal{E} : a is a *grue and green* emerald. $Ea \ \& \ (Oa \ \& \ Ga)$.
- The first part of Goodman's argument involves identifying an evidential claim that Hempel-confirms H_1 and H_2 . E_1/E_2 do not fit the bill. Why?
- As Goodman points out (more detail later), \mathcal{E} Hempel-confirms both H_1 and H_2 . Goodman thinks this is "bad news" for Hempel's theory. *Why?*