The Paradox of Confirmation

Branden Fitelson
University of California–Berkeley

- Overview of the Talk
  - The Original Formulation of the Problem
  - The Responses of Hempel/Goodman
  - Quine's Response
  - Bayesian-Inspired Clarifications of the Paradox
  - Bayesian Responses
  - A (Small) New Bayesian Result & Approach (w/ Jim Hawthorne)
  - Time Permitting: Toward an Analogous Treatment of Grue

The Paradox of the Ravens: The Original Formulation

- Nicod Condition (NC): For any object x and any properties F and G, the proposition that x has both F and G confirms the proposition that every F has G. This is a second order condition:
  $(\forall F)(\forall G)(\forall x)[Fx \,\&\, Gx$ confirms $(\forall x)(Fx \supset Gx)]$
- Equivalence Condition (EC): For any propositions H_1, E, and H_2, if E confirms H_1 and H_1 is (classically) logically equivalent to H_2, then E confirms H_2. This is also a (weak) second order condition:
  $(\forall E)(\forall H_1)(\forall H_2)[E$ confirms H_1 and $H_1 \dashv\vDash H_2 \Rightarrow E$ confirms $H_2]$
- Paradoxical Conclusion (PC): The proposition that a is both nonblack and a nonraven confirms the proposition that every raven is black. $[\sim Ba \,\&\, \sim Ra$ confirms $(\forall x)(Rx \supset Bx)]$

✓ (1) By (NC), $\sim Ba \,\&\, \sim Ra$ confirms $(\forall x)(\sim Bx \supset \sim Rx)$.
(2) By Logic, $(\forall x)(\sim Bx \supset \sim Rx) \dashv\vDash (\forall x)(Rx \supset Bx)$.
∴ (3) By (1), (2), and (EC), $\sim Ba \,\&\, \sim Ra$ confirms $(\forall x)(Rx \supset Bx)$.

Hempel's (and Goodman's) Response to the Original Paradox

- Hempel and Goodman *embraced* (NC), (EC) *and* (PC). They saw **no paradox** here. Here is Hempel's explanation (Goodman's is similar):

  > ...in the seemingly paradoxical cases of confirmation, we are often not judging the relation of the given evidence E alone to the hypothesis H ...instead, we tacitly introduce a comparison of H with a body of evidence which consists of E in conjunction with an additional amount of information we happen to have at our disposal. If the evidence E consists only of one object which ...is black [Ba], then E may reasonably be said to confirm that all objects are black [$(\forall x)Bx$], and *a fortiori* E supports the weaker assertion that all ravens are black [$(\forall x)(Rx \supset Bx)$]. [they tell a similar story for $\sim Ra$]

- H & G presuppose the *Special Consequence Condition* (SCC) here. (SCC) $(\forall E)(\forall H_1)(\forall H_2)[E$ confirms H_1 and $H_1 \vDash H_2 \Rightarrow E$ confirms $H_2]$. Contemporary Bayesians reject (SCC). More on this dialectic below.

Quine's Response to the Original Paradox

- Quine rejects (PC) but accepts (EC). ∴ He rejects (NC), *i.e.*, step (1). He argues that $\forall F$ and $\forall G$ in (NC) must be *restricted in scope*:

  (NC′) $(\forall F' \in \mathbf{N})(\forall G' \in \mathbf{N})(\forall x)[F'x \,\&\, G'x$ confirms $(\forall x)(F'x \supset G'x)]$

  He calls properties F', G' satisfying (NC′) *natural kinds*. (NC′) is a distinguishing feature of natural kinds – often called *projectibility*.
- Many (*e.g.*, H & G) are inclined to follow Quine in restricting (NC) to "natural kinds" (*e.g.*, "GRUE"). But, many (*e.g.*, H & G) reject Quine's classification of $\sim R$ and $\sim B$ *in particular* as "unnatural".
- Quine himself thinks R and B *are* "natural" (hence "projectible"). This may seem odd, but there is a history [in metaphysics] of denying the "naturalness" of "negative properties" (*i.e.*, denials of "naturals").
- Armstrong, Kim *et al* argue that "negative properties" (and other "non-naturals") can't participate in *causal relations* or *laws*. Must *this* preclude participation in *confirmation relations*? More, below.

Bayesian Confirmation Theory & The Paradox 1: Three Types of Approach

- 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 rejcted too].
  - **Comparative**. Argue that $\mathfrak{c}(H, E_1 \mid K_\alpha) > \mathfrak{c}(H, E_2 \mid K_\alpha)$, for our *actual* background knowledge K_α. Traditionally, these approaches *accept* (PC) and do *not* deny (NC) — they *entail* $\mathfrak{c}(H, E_1 \mid K_\alpha) > \mathfrak{c}(H, E_2 \mid K_\alpha) > 0$.
  - **Quantitative**. Argue that $\mathfrak{c}(H, E_2 \mid K_\alpha)$ 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 $\mathfrak{c}(H, E_1 \mid K_\alpha) \gg \mathfrak{c}(H, E_2 \mid K_\alpha) > 0$ *and* that $\mathfrak{c}(H, E_2 \mid K_\alpha) \approx 0$.
- Next, I'll critically discuss the tradition, and then describe a new approach.

Bayesian Confirmation Theory & The Paradox 2: Qualitative Approaches 1

- Most qualitative Bayesian approaches begin by making (NC) [and (PC)] more precise. Since Bayesian confirmation theory says that confirmation is a *three-place* relation [$C(H, E \mid K)$], we need some *quantifier* over the implicit K's in the traditional formulation of (NC). Here are 4 renditions:

$$(\mathrm{NC}_w) \quad (\exists K)(\forall F)(\forall G)(\forall x)[C((\forall x)(Fx \supset Gx), Fx \,\&\, Gx \mid K)]$$
$$(\mathrm{NC}_\alpha) \quad (\forall F)(\forall G)(\forall x)[C((\forall x)(Fx \supset Gx), Fx \,\&\, Gx \mid K_\alpha)]$$
$$(\mathrm{NC}_\top) \quad (\forall F)(\forall G)(\forall x)[C((\forall x)(Fx \supset Gx), Fx \,\&\, Gx \mid K_\top)]$$
$$(\mathrm{NC}_s) \quad (\forall K)(\forall F)(\forall G)(\forall x)[C((\forall x)(Fx \supset Gx), Fx \,\&\, Gx \mid K)]$$

- (NC_w) is the weakest precisification of (NC), which asserts merely that (NC) holds relative to *some* background knowledge K [(NC_w) is *too weak*].
- (NC_s) says (NC) holds for *arbitrary* background K [(NC_s) *too strong*].
- (NC_α) and ($\mathrm{NC}_\top$) are stronger than (NC_w), but weaker than (NC_s). (NC_α) says (NC) holds relative to our *actual* background knowledge K_α. And, ($\mathrm{NC}_\top$) says (NC) holds relative to "tautological" (or "empirically vacuous") background knowledge $K_\top$. These will be the salient renditions of (NC).

Bayesian Confirmation Theory & The Paradox 3: Qualitative Approaches 2

- I.J. Good was one of the first to show that the strong (Bayesian) rendition (NC_s) of Nicod's condition is false. He gave this counterexample:
  - K: Exactly one of the following hypotheses is true: (H) there are 100 black ravens, no nonblack ravens, and 1 million other birds, or else ($\sim H$) there are 1,000 black ravens, 1 white raven, and 1 million other birds. And, we are sampling at random from the universe (making all the standard probabilistic assumptions about random sampling).
  - E: a bird a is selected at random from all the birds, and Ra & Ba.
  - So, H asserts that $(\forall x)(Rx \supset Bx)$, and E is a positive instance Ra & Ba of H. With these assumptions about K, E, H, and $\sim H$, we have:

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

  - Therefore, (NC_s) is false, and *even for "natural kinds"* (*pace* Quine). Similar examples can be generated to show that (PC_s) is also false.
- So? Hempel seems to want ($\mathrm{NC}_\top$) *not* (NC_s) – see Maher (& more below).

Bayesian Confirmation Theory & The Paradox 4: Qualitative Approaches 3

- I.J. Good also claimed to have a counterexample to the more salient ($\mathrm{NC}_\top$). This one is rather controversial, however. Here's what Good says:

  > ...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 ($\mathrm{NC}_\top$). Maher gives good reason to doubt this is a counterexample to ($\mathrm{NC}_\top$).
- However, Maher more recently provides a very compelling (Carnapian) counterexample to ($\mathrm{NC}_\top$), which is beyond our scope (but see below).
- Qualitative Bayesians need to refute ($\mathrm{NC}_\top$) [or at least (NC_α) – see below]. (NC_s) is false, but Hempel knew that, and (NC_w) is too weak to do the job.

Bayesian Confirmation Theory & The Paradox 5: Comparative & Quantitative 1

- There have been *many* comparative Bayesian approaches (see Earman and Vranas for references). Here is a recent, canonical, comparative approach. Assume that our *actual* background knowledge K_α justifies the following:
  1. $\Pr(\sim Ba \mid K_\alpha) > \Pr(Ra \mid K_\alpha)$
  2. $\Pr(Ra \mid H \,\&\, K_\alpha) = \Pr(Ra \mid K_\alpha)$
  3. $\Pr(\sim Ba \mid H \,\&\, K_\alpha) = \Pr(\sim Ba \mid K_\alpha)$[a]
- **Theorem**. Any probability function satisfying (1) and (2) and (3) will also be such that (4) $\Pr(H \mid Ra \,\&\, Ba \,\&\, K_\alpha) > \Pr(H \mid \sim Ba \,\&\, \sim Ra \,\&\, K_\alpha)$.
- That is, if (1) there are (proportionally) fewer ravens than non-black things, and (2)/(3) whether something a (sampled at random from the universe) is a raven/black is independent of whether all ravens are black [H], then (4) $Ra \,\&\, Ba$ confers a higher probability on H than $\sim Ba \,\&\, \sim Ra$ does.

[a] Sometimes, a two-stage sampling model is used in which two objects a and b are sampled, and $K_\alpha \models Ra \,\&\, \sim Bb$. This yields (2), (3′) $\Pr(\sim Bb \mid H \,\&\, K_\alpha) = \Pr(\sim Bb \mid K_\alpha)$, and (4′) $\Pr(H \mid Ra \,\&\, Ba \,\&\, K_\alpha) > \Pr(H \mid \sim Bb \,\&\, \sim Rb \,\&\, K_\alpha)$. We have no real loss of generality here. See Vranas *fn*. 10.

Bayesian Confirmation Theory & The Paradox 6: Comparative & Quantitative 2

- In fact, (1)–(3) imply much more that just (4). They 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. $\mathfrak{c}(H, Ra \,\&\, Ba \mid K_\alpha) > \mathfrak{c}(H, \sim Ba \,\&\, \sim Ra \mid K_\alpha)$, for all 4 measures d, r, l, s.
- In other words, the canonical *comparative* Bayesian assumptions entail the *qualitative* claims that each of $Ra \,\&\, Ba$ and (PC) $\sim Ba \,\&\, \sim Ra$ confirms H.
- So, the canonical comparative approach is *inconsistent* with denying the salient instances of (NC_α) – it fails to isolate assumptions that undergird a comparative approach that is also consistent with a qualitative$_\alpha$ approach.
- A canonical *quantitative* approach is obtained just 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* $\Pr(H \mid \sim Ba \,\&\, \sim Ra \,\&\, K_\alpha) \approx \Pr(H \mid K_\alpha)$. That's the traditional comp/quant story.

Bayesian Confirmation Theory & The Paradox 7: Comparative & Quantitative 3

- 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. This is not so controversial.
- The controversial assumptions are the independence assumptions (2) and (3). Recently, Vranas has provided compelling reasons to worry about (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 Bayesian approaches (done right). First, note that we can replace (2) and (3) with the far weaker:

  (‡) $\Pr(H \mid Ra \,\&\, K_\alpha) \geqslant \Pr(H \mid \sim Ba \,\&\, K_\alpha)$
- (1) and (‡) $\Rightarrow \Pr(H \mid Ra \,\&\, Ba \,\&\, K_\alpha) > \Pr(H \mid \sim Ba \,\&\, \sim Ra \,\&\, K_\alpha)$.
- And, (1) and (‡) are consistent with denying the salient instances of (NC_α), *i.e.*, with both $Ra \,\&\, Ba$ and $\sim Ba \,\&\, \sim Ra$ *dis*confirming H (rel. to K_α)!

Bayesian Confirmation Theory & The Paradox 8: Comparative & Quantitative 4

- Here's is an interesting fact about the three measures d, r, and l:

  (i) For all H, E_1, E_2, and K, and for measures $\mathfrak{c} = d$, $\mathfrak{c} = r$, and $\mathfrak{c} = l$:
  $\Pr(H \mid E_1 \,\&\, K) > \Pr(H \mid E_2 \,\&\, K) \Rightarrow \mathfrak{c}(H, E_1 \mid K) > \mathfrak{c}(H, E_2 \mid K)$.
- Thus, we know that (1) and (†) must be jointly sufficient for the desired inequality $\mathfrak{c}(H, E_1 \mid K) > \mathfrak{c}(H, E_2 \mid K)$, for each of the three measures $\mathfrak{c} = d$, $\mathfrak{c} = r$, and $\mathfrak{c} = l$. But, what about $\mathfrak{c} = s$? Surprisingly, the answer is no!

  $$(1) \text{ and } (‡) \not\Rightarrow s(H, Ra \,\&\, Ba \mid K) > s(H, \sim Ba \,\&\, \sim Ra \mid K)$$
- In fact, (1) and (‡) $\not\Rightarrow s(H, Ra \,\&\, Ba \mid K) > s(H, \sim Ba \,\&\, \sim Ra \mid K)$ *even if* we add the assumption that both $Ra \,\&\, Ba$ and $\sim Ba \,\&\, \sim Ra$ confirm H!
- The fact that s says a non-black non-raven can confirm that all ravens are black more strongly than a black raven does is bad news for s. In fact, the violation of (i) by s has caused one of its early defenders to abandon it!
- This makes two senses in which the standard approach is too strong. It conflates the qualitative, comparative, and quantitative, and it obscures an undesirable property of s. Hawthorne & I have more general results …

Bayesian Confirmation Theory & The Paradox 9: Nicod's Condition Again

- As I have shown, most Bayesians have given comparative "resolutions" of the paradox that commit them to *accepting* (PC). As a result, the standard resolutions are inconsistent with denying (NC_α) *in this instance*.
- Our account allows the Bayesian to reject (NC_α) even in this very case, and still to maintain the comparative claim they desire. The question still remains, however, whether (NC) in its salient form should be accepted.
- This depends on whether we are doing *epistemology* or *logic*. If we are doing epistemology, then (NC_α) seems the salient rendition. And, I claim, it is unclear, given K_α (*viz.*, what we know), whether E_1 or E_2 confirms H.
- What seems clear is that E_1 confirms H more strongly (or disconfirms it less strongly) than E_2 does. This claim rests only on very weak assumptions about K_α that I think most people (even Vranas!) would find acceptable.
- It is unclear whether the stronger assumptions in the standard "resolutions" are satisfied, given what we actually know about the universe (see Vranas).

- Besides, epistemologically, the question is not whether (NC_α) is true *in this instance*, but whether it is true *generally* (*i.e.*, for *all* F, G, and a). I think it is pretty clear that (NC_α) is false for *some* F, G, and a. The question is *why*.
- Here, I think Quine's answer is unhelpful (and without theoretical merit). He tries to locate the failure of (NC_α) in the "naturalness" of F and G, thereby restricting the scope of the second-order quantifiers over F and G in the hopes of cordoning off the "true domain of application" of (NC_α).
- A more theoretically rooted answer might be (using the modern Bayesian machinery) that (NC_α) is true if $\Pr(H \mid Fa \,\&\, Ga \,\&\, K_\alpha) > \Pr(H \mid K_\alpha)$. Perhaps Quine thinks that only "natural kinds" can satisfy this inequality.
- But, surely, this inequality often holds for "non-natural" properties F, G. All it says is that $Fa \,\&\, Ga$ is *correlated* with H under some *epistemically permissible* credence function. Of course, this correlation is *fallible*.
- Rational credence functions are bound sometimes to (permissibly) reflect "spurious" correlations (*i.e.*, correlations that are not grounded in causal or lawlike structure). Do no *evidential* relations exist in such cases? VENN:

> The science of Probability makes no assumption whatever about the way in which events are brought about, whether by causation or without it. All that we undertake to do is to establish . . . a body of rules which are applicable to classes of cases in which we do not or cannot make inferences about the individuals.

- From a *logical* point of view, it is unclear which rendition of (NC) is salient. If you think that in order to be "logical", inductive logic has to provide relations between propositions relative to some "logical aether", then you'll be inclined to think that ($\mathrm{NC}_\top$) is the salient instance of (NC).
- This is what Hempel thought. This is why he abandoned the probabilistic inductive logic project (he thought no requisite account of "logical probability" was forthcoming), and he moved to his "instantial" account.
- On this score, Hempel's (and Goodman's) argument for the truth of ($\mathrm{NC}_\top$) is laden with suspect theoretical commitments [in particular, (SCC)].
- Since contemporary Bayesians reject (SCC), Hempel and Goodman have given them no compelling reason to think that ($\mathrm{NC}_\top$) is true. And, Maher's recent work seems to indicate that ($\mathrm{NC}_\top$) is false. Besides, who cares?

- Why should we think ($\mathrm{NC}_\top$) is *salient* here? As I just said, the only reason one would be inclined to think that ($\mathrm{NC}_\top$) is "*the* logically salient rendition of (NC)" would be if one thought that "relativity to $\top$" is what makes inductive judgments *logical*. But, I reject that *and* Carnap's programme.
- The picture of inductive logic as a "limiting case of epistemology" where our background knowledge "tends to zero" is wrongheaded. And, so is the tendency to think that ($\mathrm{NC}_\top$) has any *special* logical relevance here.
- From the point of view of inductive logic, any probability model is as good as any other. There is no "logically privileged" model. Proper "logical" accounts of probability and/or confirmation make no such commitments.
- When we *apply* inductive logic (or probability), we need to ask what is the right model to use for the purpose at hand. *That*, I submit, is not to be decided by logic, but, by epistemology (or rational decision theory).
- If so, then we're back to (NC_α) again, or so it would seem, from a Bayesian point of view. But, I think the real question is which $\mathcal{M}$ is the right one for modeling the evidential relation between E_1, E_2, and H. Food for thought.

Goodman's "Grue" Paradox: The Standard Formulation

- Goodman presents (and Hempel discusses) an example involving the following two hypotheses (H and H') and observation report ($\mathcal{E}$):

  H: All emeralds are green. $[(\forall x)(Ex \supset Gx)]$

  H': All emeralds are grue. $[(\forall x)(Ex \supset \mathcal{G}x)]$

  $\mathcal{E}$: An object a has been observed to be a green emerald $[Ea \,\&\, Ga]$.

- The predicate "grue" is defined as follows [note what this *doesn't* say]:

  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 not green.

- Thus, $\mathcal{E}$ is equivalent to $Ea \,\&\, \mathcal{G}a$, and so $\mathcal{E}$ is a "positive instance" of both H and H'. So, by (NC), $\mathcal{E}$ confirms both H and H' – seems paradoxical.
- The same sort of Bayesian approaches we saw for the Ravens have been tried for GRUE. One can reject (NC), or one can try to argue that H is better confirmed by $\mathcal{E}$ than H' is, relative to our *actual* K_α. That is, $\mathfrak{c}(H, \mathcal{E} \mid K_\alpha) > \mathfrak{c}(H', \mathcal{E} \mid K_\alpha)$. But, first, some non-Bayesian approaches.

Goodman's "Grue" Paradox: Non-Bayesian Approaches

- Hempel and Goodman (and Carnap, and many, many others) change their tune about appealing to "naturalness" of the predicate $\mathcal{G}$ in this case.
- Most non-Bayesians who write about this paradox think that the way to resolve it is to restrict the scope of the G quantifier in (NC) to "naturals".
- We have already seen that this is unnecessary, since (NC) [and (SCC)] is false even for "natural kinds", and so the paradox doesn't get off the ground. But, I'll say a few things before looking at Bayesian approaches.
- The idea seems to be that there is something "gerrymandered" or "unnatural" about $\mathcal{G}$. But, it is quite difficult to say what this means. And, it is even more difficult to say why *that* should matter for *confirmation*.
- Goodman's own "solution" was to claim that we are willing to project predicates because they are *entrenched* in our language and previous practices of reasoning about and interacting with the world.

- But this is really no solution at all, rather the acknowledgement of certain prejudices. It is also open to the objection of being unduly conservative.
- Commonsense predicates seem better entrenched than their scientific replacements, but we don't think of this as a good reason for preferring commonsense descriptions to scientific ones for purposes of prediction.
- Some (*e.g.*, Carnap) have tried to argue that grue is not symmetrical with respect to confirmation because: (1) it isn't purely qualitative, but "positional" – it involves reference to a particular time, or (2) it predicts a change, from green to blue, and we have no evidence suggesting a change.
- Goodman: what if someone speaks of and thinks in terms of grue (rather than green)? From that person's point of view it's our predicate green that is "positional" (for a thing is green just in case it is grue up to t and non-grue thereafter). So, this seems no better than "entrenchment".
- There are *many* non-probabilistic approaches to GRUE, but I will not discuss them, since I think they are unnecessary, and our focus is on Pr.
- So, we now move on to comparative Bayesian resolutions of the paradox.

Eells on Goodman's "Grue" Paradox

- Eells offers a Bayesian account of the Grue paradox which trades on the following property of d, where β and δ are defined as follows:
  - $\beta =_{df} \Pr(H \,\&\, E) - \Pr(H' \,\&\, E)$, and
  - $\delta =_{df} \Pr(H \,\&\, {\sim}E) - \Pr(H' \,\&\, {\sim}E)$

  (1) If $\beta > \delta$ and $\Pr(E) < \dfrac{1}{2}$, then $d(H,E) > d(H',E)$.

- Problem: Neither r nor l satisfies property (1).
- Eells does provide reasons (as reported in a paper by Sober, see below) to prefer the difference measure d over the ratio measure r, but he does not supply any reasons to prefer d over l.
- Moreover, I have (earlier lecture) provided reasons to favor l *over* d.
- Also, why think the antecedent of (1) holds in the GRUE case?

Sober on Goodman's "Grue" Paradox

- Sober gives a more robust Bayesian account, which uses the following:

(2) If H, H' entail E, then $d(H, E) > d(H', E)$ iff $\Pr(H) > \Pr(H')$.

- r violates even this weaker condition, but l *does* satisfy (2).
- So, Sober's resolution of Goodman's "Grue" paradox is *less* sensitive to choice of measure than Eells's is. And, since (2) holds for l, it's Kosher.
- But, it is unclear why (2) is *relevant* here, since the antecedent of (2) is not satisfied in the Grue paradox! Sober seems to be assuming that H_1 and H_2 each entail E. But, in fact, *neither* entails E. Why is that?
- Moreover, Sober's resolution of Grue saddles the Bayesian with the view that the evidence cannot, *a posteriori*, favor H over H' (or *vice versa*). If Sober is right, then only their relative *a priori* plausibility is relevant.
- Note how different this is from Bayesian resolutions of the Ravens. What we want is an *a posteriori* way to adjudicate between H and H', using E.

Goodman's "Grue" Paradox: Setting-Up a Finer-Grained Representation

- To capture the full logical structure of Goodman's paradox (as we did with Hempel's), we need three $\forall$-hypotheses and three predicates:

H_1: All observed emeralds are green. $[(\forall x)(Ex \supset (Ox \supset Gx))]$

H_2: All unobserved emeralds are green. $[(\forall x)(Ex \supset (\sim Ox \supset Gx))]$

H_3: No unobserved emeralds are green. $[(\forall x)(Ex \supset (\sim Ox \supset \sim Gx))]$

$\mathcal{E}$: $Ea \,\&\, Oa \,\&\, Ga$ [for some a sampled from the universe]

- Now, we can see that $H = H_1 \,\&\, H_2$, and $H' = H_1 \,\&\, H_3$. And, what we want to know is how strongly $\mathcal{E}$ differentially confirms H and H', relative to (say) our actual background knowledge (K_α). This is non-trivial!
- Jim Hawthorne and I are hard at work on this. You can see why Bayesians have ignored the logical fine-structure here! If you don't, you've got a very nasty set of probability problems to crunch!
- We don't know much about the full story yet. But, the goal is to use assumptions similar to those used in the Ravens Paradox, *e.g.*, $\Pr(\sim Ga \mid K_\alpha) \gg \Pr(Ea \mid K_\alpha)$, and $\Pr(\sim Oa \mid K_\alpha) \gg \Pr(Ea \mid K_\alpha)$, *etc.*