Philosophy 148 — Announcements & Such

- Raul is away for a few days this week. Three Announcements:
  - The make up section for Raul's regular Tuesday 10–11 section will be on Thursday 10–11 @ 206 Wheeler.
  - The make up section for Raul's regular Wednesday 9–10 section will be on Friday 12–1 (location TBA on the website and in class Thurs).
  - Raul's make-up OH are Friday 11–12 & 1–2 @ 5323 Tolman.
- HW #5 has been posted — it's due next Thursday.
  - It's on Hempel's theory of confirmation (some important formal properties), and a Bayesian theorem pertaining to the raven paradox.
- Today's Agenda
  - The Raven Paradox
  - Next: The Grue Paradox
  - Final Topic: Psychological Applications of Confirmation Theory

The Raven Paradox (aka., The Paradox of Confirmation)

- **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. Strong second-order condition:

  $$(\forall F)(\forall G)(\forall x)[Fx \,\&\, Gx \text{ confirms } (\forall x)(Fx \supset Gx)]$$

- **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. Weak 2^{nd} order condition:

  $$(\forall E)(\forall H_1)(\forall H_2)[E \text{ confirms } H_1 \text{ and } H_1 \dashv\vDash H_2 \Rightarrow E \text{ confirms } H_2]$$

- **Paradoxical Conclusion** (PC): The proposition that a is both nonblack and a nonraven confirms the proposition that every raven is black. This is a first-order condition (arbitrary a): $\sim Ba \,\&\, \sim Ra$ confirms $(\forall x)(Rx \supset Bx)$.

Proof. (1) By (NC), $\sim Ba \,\&\, \sim Ra$ confirms $(\forall x)(\sim Bx \supset \sim Rx)$.

(2) By Classical Logic, $(\forall x)(\sim Bx \supset \sim Rx) \dashv\vDash (\forall x)(Rx \supset Bx)$.

$\therefore$ (PC) By (1), (2), (EC), $\sim Ba \,\&\, \sim Ra$ confirms $(\forall x)(Rx \supset Bx)$.

- Hempel & Goodman *embraced* (NC), (EC) *and* (PC). They saw **no paradox**. Hempel *explains away* the paradoxical *appearance* (Goodman does same):

  > … 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 … E in conjunction with … additional … information we … have at our disposal.

- Hempel's Idea: E [$\sim Ra \,\&\, \sim Ba$] confirms H [$(\forall x)(Rx \supset Bx)$] *relative to* $\top$, but E doesn't confirm H relative to some (nontautological) $K \neq T$.

- *Which* $K \neq \top$? Later, Hempel discusses $K = \sim Ra$. Intuition: if you already know that a is a nonraven, then observing its color will not tell you anything about the color of ravens. Hempel: (PC) is true, but (PC*) is false:

(PC) $\sim Ra \,\&\, \sim Ba$ confirms $(\forall x)(Rx \supset Bx)$, relative to $\top$.

(PC*) $\sim Ra \,\&\, \sim Ba$ confirms $(\forall x)(Rx \supset Bx)$, relative to $\sim Ra$.

- This is a good insight! Unfortunately, it is *logically incompatible* with the (deductive) confirmation *theories* that Hempel and Goodman accept.

- Specifically, this possibility contradicts the *K-monotonicity* property:

(M_K) E confirms H, relative to $\top \Rightarrow E$ confirms H relative to *any* K (provided that K does not mention any individuals not already mentioned in E).

- Because Hempel's theory of confirmation satisfies (M), his theory implies that (PC) entails (PC*). So, it is logically impossible for Hempel's theory to undergird his suggestion that (PC) is true, while (PC*) is false.
- This is bad news for Hempel/Goodman. Surprisingly, nobody noticed this inconsistency in the Hempel/Goodman approach to the paradox.
- As we will see shortly, *Bayesians* can better accommodate Hempel's intuitions here, since *their* theories of confirmation do *not* satisfy (M).
- Interestingly, later in this very same passage, Hempel offers an argument for premise (1) which, itself, *depends on* (M)! [See my handout.]

  > If ... E consists *only* of one ... nonraven [$\sim Ra$], then E ... confirm[s] that all objects are nonravens [$(\forall x)\sim Rx$], and *a fortiori* E supports the weaker assertion that all nonblack objects are nonravens [$(\forall x)(\sim Bx \supset \sim Rx)$].

- The dependence on (M) is almost invisible here! My conjecture: (M) is a vestige of "objectual" ways of thinking about confirmation (like NC_0).

- Scheffler rejects (PC), but accepts (1). He denies (EC). He proposes:

(*) A Hempelian positive instance (E) of a $\forall$-hypothesis (H) confirms H, *unless E is also a positive instance of the contrary (H') of H*.

- Let H: $(\forall x)[Rx \supset Bx]$. The contrary of H is H': $(\forall x)[Rx \supset \sim Bx]$. Let E: $\sim Ra \,\&\, \sim Ba$. E is a Hempelian positive instance of H, *and* H'.

- Thus, according to Scheffler's (*), E does not confirm H after all.

- Scheffler accepts (1) [and (NC)]. E confirms H^*: $(\forall x)[\sim Bx \supset \sim Rx]$ — *even according to* (*). This is because E is *not* a Hempelian positive instance of the *contrary* of H^*, $H^{*\prime}$: $(\forall x)[\sim Bx \supset Rx]$.

- This leads to a violation of (EC), of course, since – according to (*) – E confirms H^*, but E does not confirm H — even though $H \mathrel{=\!\mid\!\models} H^*$.

- Is Scheffler's (*) true? **Exercise**: show that Scheffler's (*), and (NC) are both *false* from the point of view of PR-theory. I'll return to this when we discuss I.J. Good and (NC). This will be one of the many subtle (and non-Hempelian) aspects of of probabilistic relevance accounts of $\mathfrak{c}$.

- Quine rejects (PC) but accepts (EC). As a result, Quine rejects (1), and he argues that $\forall F$ and $\forall G$ in (NC) must be *restricted in scope*:

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

- Quine calls properties F', G' satisfying (NC′) "*projectible*." And, he says that *natural kinds* are distinctively projectible in this sense.

- 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 thinks R and B *are* "natural" (hence "projectible"). This may seem odd, but there is a history [esp. in metaphysics] of denying the "naturalness" of "negative properties" (denials of "naturals").

- Some have followed Quine (Kim Quines the confirmation of psychological laws). Why should "non-naturalness" rule-out *confirmation*? And, what's Quine's *theory* of confirmation?

- This is for another course! Meanwhile, let's look at modern Bayesian confirmation theory, which is "subjective" / "epistemic".

Digression on Bayesian Confirmation Theory 1

- Contemporary Bayesian confirmation theory: E *confirms* H relative to K [$\mathfrak{C}(H, E \mid K)$] if E and H are conditionally correlated, given K [under a subjective/epistemic $\Pr_t^s(\cdot)$ function, not a "logical" one, like Carnap's].
- If E and H are conditionally anti-correlated, given K [under $\Pr_t^s(\cdot)$], then E *dis*confirms H relative to K (for s, at t), and if $E \perp\!\!\!\perp H \mid K$ [under $\Pr_t^s(\cdot)$], then E is *confirmationally neutral* regarding H relative to K (for s, at t).
- There are *many* logically equivalent (but syntactically different) ways of saying that E confirms H (relative to K). Here are three of these ways:
  - E confirms H, relative to K if $\Pr(H \mid E \,\&\, K) > \Pr(H \mid K)$.
  - E confirms H, relative to K if $\Pr(E \mid H \,\&\, K) > \Pr(E \mid {\sim}H \,\&\, K)$.
  - E confirms H, relative to K if $\Pr(H \mid E \,\&\, K) > \Pr(H \mid {\sim}E \,\&\, K)$.
- By taking differences, ratios, *etc.*, of the left/right sides of such inequalities, *many quantitative* Bayesian *relevance measures* $\mathfrak{c}(H, E \mid K)$ of the *degree* to which E confirms H (relative to K) can be constructed.

Digression on Bayesian Confirmation Theory 2

- *Dozens* of Bayesian relevance measures of the degree to which E confirms H (relative to K) have been proposed. Here are four representative choices for $\mathfrak{c}$, generated using the three inequalities on the previous slide.
  - The *Difference*: $d(H, E \mid K) = \Pr(H \mid E \,\&\, K) - \Pr(H \mid K)$
  - The *Ratio*: $r(H, E \mid K) = \dfrac{\Pr(H \mid E \,\&\, K)}{\Pr(H \mid K)} \doteq \dfrac{\Pr(H \mid E \,\&\, K) - \Pr(H \mid K)}{\Pr(H \mid E \,\&\, K) + \Pr(H \mid K)}$
  - The *Likelihood-Ratio*:
  $$l(H, E \mid K) = \frac{\Pr(E \mid H \,\&\, K)}{\Pr(E \mid {\sim}H \,\&\, K)} \doteq \frac{\Pr(E \mid H \,\&\, K) - \Pr(E \mid {\sim}H \,\&\, K)}{\Pr(E \mid H \,\&\, K) + \Pr(E \mid {\sim}H \,\&\, K)}$$
  - The *Normalized-Difference*:
  $$s(H, E \mid K) = \Pr(H \mid E \,\&\, K) - \Pr(H \mid {\sim}E \,\&\, K) = \frac{1}{\Pr({\sim}E \mid K)} \cdot d(H, E \mid K)$$
- *A fortiori*, *all* Bayesian confirmation measures agree on *qualitative* judgments like "E confirms/disconfirms/is irrelevant to H". *But*, . . .

Digression on Bayesian Confirmation Theory 3

- Two relevance measures c_1 and c_2 are *ordinally equivalent* ($c_1 \doteq c_2$) iff

$$c_1(H, E \mid K) \geqslant c_1(H', E' \mid K') \text{ iff } c_2(H, E \mid K) \geqslant c_2(H', E' \mid K')$$

- Surprisingly, none of the four most popular relevance measures is ordinally equivalent with any of the others! This has ramifications.
- If an argument $\mathcal{A}$ is valid for some choices of measure of confirmation, but invalid for other choices (holding fixed all other assumptions in the premises of the argument), then $\mathcal{A}$ is *sensitive to choice of measure*.
- For instance, it is part of Bayesian Lore that the observation of a black raven (E_1) confirms the hypothesis that all ravens are black (H) more strongly than the observation of a red herring (E_2) does. But, this conclusion *depends sensitively on one's choice of confirmation measure*.
- *Almost all* comparative arguments are *sensitive to choice of measure*!

Digression on Bayesian Confirmation Theory 4

- Very few attempts to resolve the problem of measure-sensitivity have appeared in the Bayesian literature. And, those arguments that have appeared tend to be very technical and unintuitive. We can do better.
- When we discuss Inductive Logic (after Bayesian confirmation theory), we will see that l stands out from the other relevance measures. l is the only (adequate) relevance measure that gets the deductive cases right.
- That is, l is the only (adequate) $\mathfrak{R}$-measure in the literature that satisfies:

$$\mathfrak{c}(H, E \mid K) \text{ should be } \begin{cases} \text{Maximal } (> 0, \text{ constant}) & \Leftarrow E \,\&\, K \text{ entails } H. \\ > 0 \text{ (confirmation rel. to } K) & \Rightarrow \Pr(H \mid E \,\&\, K) > \Pr(H \mid K). \\ = 0 \text{ (irrelevance rel. to } K) & \Rightarrow \Pr(H \mid E \,\&\, K) = \Pr(H \mid K). \\ < 0 \text{ (disconfirmation rel. to } K) & \Rightarrow \Pr(H \mid E \,\&\, K) < \Pr(H \mid K). \\ \text{Minimal } (< 0, \text{ constant}) & \Leftarrow E \,\&\, K \text{ entails } {\sim}H. \end{cases}$$

- Here, we assume that $\mathfrak{c}$ is *defined*, which constrains the unconditional Pr's.

Digression on Bayesian Confirmation Theory 5

- Here's how our four relevance measures handle the deductive cases.

- $l(H, E \mid K) = \begin{cases} +1 & \text{if } E \,\&\, K \vDash H, \Pr(E \mid K) > 0, \Pr(H \mid K) \in (0,1) \\ -1 & \text{if } E \,\&\, K \vDash {\sim}H, \Pr(E \mid K) > 0, \Pr(H \mid K) \in (0,1) \end{cases}$

- $d(H, E \mid K) = \begin{cases} \Pr({\sim}H \mid K) & \text{if } E \,\&\, K \vDash H, \Pr(E \mid K) > 0 \\ -\Pr(H \mid K) & \text{if } E \,\&\, K \vDash {\sim}H, \Pr(E \mid K) > 0 \end{cases}$

- $r(H, E \mid K) = \begin{cases} \dfrac{1 - \Pr(H \mid K)}{1 + \Pr(H \mid K)} & \text{if } E \,\&\, K \vDash H, \Pr(E \mid K) > 0, \Pr(H \mid K) > 0 \\ -1 & \text{if } E \,\&\, K \vDash {\sim}H, \Pr(E \mid K) > 0, \Pr(H \mid K) > 0 \end{cases}$

- $s(H, E \mid K) = \begin{cases} \Pr({\sim}H \mid {\sim}E \,\&\, K) & \text{if } E \,\&\, K \vDash H, \Pr(E \mid K) \in (0,1) \\ -\Pr(H \mid {\sim}E \,\&\, K) & \text{if } E \,\&\, K \vDash {\sim}H, \Pr(E \mid K) \in (0,1) \end{cases}$

- From an inductive-logical point of view, this seems to favor l over the other popular measures. But, what about an *epistemic* point of view? [We must be careful here, if we think of $\mathfrak{c}$ as a measure of *evidential support*!]

Digression on Bayesian Confirmation Theory 6

- Consider the following two propositions concerning a card c, drawn at random from a standard deck of playing cards (classical model $\mathcal{M}$):

E: c is the ace of spades. H: c is *some* spade.

- I take it as intuitively clear and uncontroversial that ($K = \top$ is omitted):

(S_1) The degree to which E supports $H \neq$ the degree to which H supports E, since $E \vDash H$, but $H \nvDash E$. Intuitively, we have $\mathfrak{c}(H, E) \gg \mathfrak{c}(E, H)$.

(S_2) The degree to which E confirms $H \neq$ the degree to which $\sim E$ *dis*confirms H, since $E \vDash H$, but $\sim E \nvDash \sim H$. Intuitively, $\mathfrak{c}(H, E) \gg -\mathfrak{c}(H, \sim E)$.

- Therefore, *no adequate relevance measure of support* $\mathfrak{c}$ *should be such that either* $\mathfrak{c}(H, E) = -\mathfrak{c}(H, \sim E)$ *or* $\mathfrak{c}(H, E) = \mathfrak{c}(E, H)$ (for all E and H and all models $\mathcal{M}$). I'll call these two symmetry desiderata S_1 and S_2, respectively.

- Note: $r(H, E) = r(E, H)$ and $s(H, E) = -s(H, \sim E)$. So, r violates S_1 and s violates S_2. d and l satisfy these desiderata. [This is interesting, *if* such symmetry desiderata hold for measures of *evidential support*.]

Digression on Bayesian Confirmation Theory 7

- As we'll soon see, Bayesian accounts of the Ravens Paradox presuppose that the following property holds for relevance measures $\mathfrak{c}$ that they use:

(†) If $\Pr(H \mid E_1) > \Pr(H \mid E_2)$, then $\mathfrak{c}(H, E_1) > \mathfrak{c}(H, E_2)$.

- As it turns out, not all relevance measures satisfy (†). This is quite a surprising fact. Specifically, s *violates* (†). For this reason, some former advocates of s (*e.g.*, Joyce) have now abandoned it in this context.
- Typically, as we'll see, the advocates of such arguments have used either d or r in their arguments. And, as it turns out, d, r, and l *all* satisfy (†). Indeed, (†) is almost universally accepted as a *desideratum* for $\mathfrak{c}$.
- None of these authors seems to provide (*independent*) reasons to prefer their measures over s, or other measures which violate (†). We've seen that s can be ruled-out on inductive-logical grounds or epistemic grounds.
- I will just assume that (†) is in force in my discussion of the raven paradox. But, this issue about (†) will arise again in that context.

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 un*intuitive (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 Hs).
  - Hempel: $E \,\&\, K \vDash dev_{I(E)}(H)$ [$I(E)$ = individuals mentioned in E].
  - H–D: $H \,\&\, K \vDash E$.
  - Absolute Bayes (confirms$_f$): $\Pr(H \mid E \,\&\, K) > r$, for some "threshold" r.
  - Incremental Bayes (confirms$_i$): $\Pr(H \mid E \,\&\, K) > \Pr(H \mid K)$.
- 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: $\Pr(H \mid E \,\&\, K) > \Pr(H' \mid E' \,\&\, K')$.
  - Incremental Bayes: $\mathfrak{c}(H, E \mid K) > \mathfrak{c}(H', E' \mid K')$ [some $\mathfrak{R}$-measure $\mathfrak{c}$].
- E confirms H relative to K to degree α:
  - Nicod, Hempel, H–D: No accounts (purely qualitative theories).
  - Absolute Bayes: $f(\Pr(H \mid E \,\&\, K), r) = \alpha$ [some function f, thresh. r]
  - Incremental Bayes: $\mathfrak{c}(H, E \mid K) = \alpha$ [some $\mathfrak{R}$-measure $\mathfrak{c}$].

- 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 discuss the tradition, and then describe a new approach.
- All Bayesian approaches begin by *precisifying* (NC) [and (PC)] . . .

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

(NC$_w$) $(\exists K)(\forall F)(\forall G)(\forall x)[\mathfrak{C}((\forall x)(Fx \supset Gx), Fx \,\&\, Gx \mid K)]$

(NC$_\alpha$) $(\forall F)(\forall G)(\forall x)[\mathfrak{C}((\forall x)(Fx \supset Gx), Fx \,\&\, Gx \mid K_\alpha)]$

(NC$_\top$) $(\forall F)(\forall G)(\forall x)[\mathfrak{C}((\forall x)(Fx \supset Gx), Fx \,\&\, Gx \mid K_\top)]$

(NC$_s$) $(\forall K)(\forall F)(\forall G)(\forall x)[\mathfrak{C}((\forall x)(Fx \supset Gx), Fx \,\&\, Gx \mid K)]$

- (NC$_w$) is *too weak* [let K = "all instances confirm all generalizations"].
- Hempel's "explaning away" *suggests* (NC$_s$) should be *too strong*.
- So (NC$_\alpha$) and (NC$_\top$) 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$_s$).

- 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:

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

- 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"...

- 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* "$\Pr_{\mathsf{T}}$." They've nothing to say about (NC_{T}). This is why they *abandon qualitative* approaches in favor of the *comparative/quantiative*.

- 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) $\Pr(\sim Ba \mid K_\alpha) > \Pr(Ra \mid K_\alpha)$

(2) $\Pr(Ra \mid H \,\&\, K_\alpha) = \Pr(Ra \mid K_\alpha)$ $[\therefore \sim Ra \perp\!\!\!\perp H \mid K_\alpha \text{ (!)}]$

(3) $\Pr(\sim Ba \mid H \,\&\, K_\alpha) = \Pr(\sim Ba \mid K_\alpha)$ $[\therefore Ba \perp\!\!\!\perp H \mid K_\alpha \text{ (!)}]$

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

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

- 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) …