Jeffrey Conditionalization

Phil 735/fall 2023 – Oct 25

Start with a probability distribution that we'll call Old(-), for example

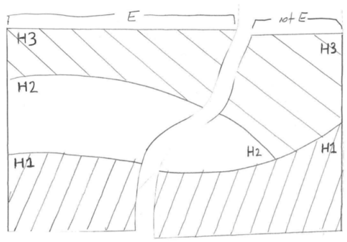

Let's make a photocopy of that, magnified by some factor which is allowed to be smaller or larger than 1.0:

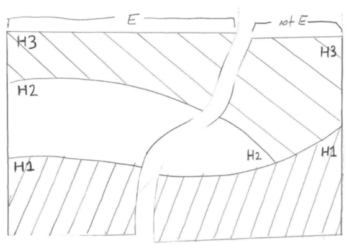

magnified by β = 1.2:

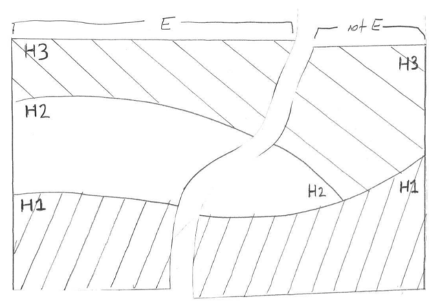

Then discard the not-E part of the photocopy, retaining the E part and joining it with the original distribution:

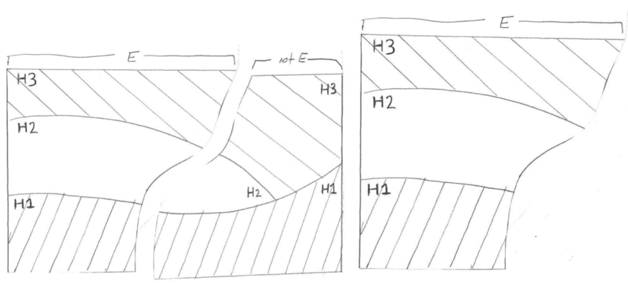

There are three regions here:

(a) the E region of the original distribution, whose size is Old(E)
(b) the not-E region of the original distribution, whose size is Old(not-E)
(c) the photocopied E region, whose size is β Old(E)

Add their sizes together and you get n = 1 + β Old(E). We can scale/renormalize all three regions by dividing them by n, so that our result adds up again to size 1. Call the result our new probability function New(-).

New(H3) = (Old(H3 & E) + Old(H3 & not-E) + Old(H3 & E) * β) / n = (Old(H3) + Old(H3 & E) * β) / n
New(E) = (Old(E) + Old(E) * β) / n

Notice that because the newly added region is a (magnified by β) copy of the original Old(E) region, the value of New(H|E) for any H will still be Old(H|E).

Another way to think of the operation we described is that, instead of adjoining a copy of the original E region, we keep a single picture, but the picture has sand or mud piled on it. We start with 1 liter of sand. Then we add β Old(E) new liters of sand, distributed only over the E region, and distributed in such a way as to maintain all the ratios of the subregions of E to each other. The new volume of sand is now n = 1 + β Old(E).

The "**X: Y Bayes Factor**" of a change from Old(-) to New(-) is defined as: (New(X)/New(Y)) / (Old(X) / Old(Y)).

Let's calculate the E: not-E Bayes Factor of the operation we've described:

$$\frac{[\text{ New(E)} = (\text{ Old(E)} + \text{Old(E)} * \beta \,) / n \,] \,/\, [\text{ New(not-E)} = \text{Old(not-E)} / n \,]}{\text{Old(E)} / \text{Old(not-E)}} = \frac{(\text{ Old(E)} * (1+\beta)\,) / \text{Old(not-E)}}{\text{Old(E)} / \text{Old(not-E)}} = \mathbf{1 + \beta}$$

Field's α = ½ log(1 + my β).
My $\beta = (e^{2\alpha}) - 1$
Since (log x) + (log y) = log xy, where I'm multiplying βs, Field will add αs. That's the main point of his using logs.

Adding a Second Operation
This time let's start with New(-) and add a new magnification factor β', ending with New´(-). For simplicity, we'll again use an E versus not-E partition, keeping the E region of the copy. But you could also consider different operations, where it's the not-E region that you keep instead, or (E $\vee$ H1) or (E $\wedge$ (H2 $\vee$ H3 $\supset$ H3)).

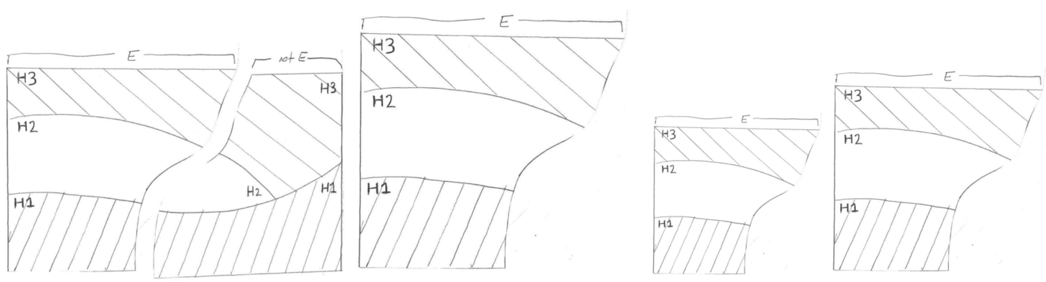

The first three regions are from New(-), and have a combined size of 1:
(a) the first region's size is Old(E) / n
(b) the second region's size is Old(not-E) / n
(c) the third region's size is Old(E) β / n

The fourth and fifth regions are copies of the E regions from New(-), and are magnified by our new factor β'. Here $\beta' < 1$.
(d) The fourth region's size is Old(E) / n * β'
(e) The fifth region's size is Old(E) β /n * β'

The combined size of regions (d) and (e) is New(E) β'.

The combined size of all five regions is 1 + New(E) β' = 1 + Old(E) (1+β) / n * β'. Let this be n´, and we scale/renormalize by dividing each region's size by n´, so that our result adds up again to size 1.
New´(H3) = (New(H3) + New(H3 & E) * β') / n´
New´(E) = (New(E) + New(E) * β') / n´

The E: not-E Bayes Factor of the change from New(-) to New´(-) will be 1 + β'.
The E: not-E Bayes Factor of the cumulative change from Old(-) to New´(-) will be:

$$\frac{\text{New´(E)} \;/\; \text{New´(not-E)}}{\text{Old(E)} / \text{Old(not-E)}} = \frac{(\text{ New(E)} * (1+\beta')\,) / \text{New(not-E)}}{\text{Old(E)} / \text{Old(not-E)}} = \mathbf{(1 + \beta)\,(1 + \beta')}$$

What would be the result of "doing these operations in reverse order"?

We want to formally represent "receiving/processing the same information/evidence in reverse order." But if it's to be *the same information*, then it shouldn't include facts about which information was received first. It shouldn't be that in one case we're thinking "she laughed then it became a scowl" and in the other case we're thinking "she scowled then it became a laugh" (Lange 2000). Our information needs to say the same thing in both cases. It'd be OK if our information was silent about the order in which it was received; or if it said it was received simultaneously; or if in one case it misrepresented which information was received first (so long as our formalism was equipped to handle false information).

If we're going to do the same operations in reverse order, we need first to *perform the same operation* that took us from New(-) to New´(-), only now we start with distribution Old(-) instead of distribution New(-). But *what counts as* "performing the same operation" on a different starting distribution? This is a substantive question, that could be answered in different ways. There isn't a single obvious choice.

We'll assume that "performing the same operation" means using the same partition (E versus not-E), and that the transformation is "rigid" with respect to the cells of that partition. Let's call this a transformation from Old(-) to New´´(-), and what this claim about "rigidity" means is that for every H, New´´(H|E) = Old(H|E) and New´´(H|not-E) = Old(H|not-E).

That still leaves further choices unsettled.

- One idea for "applying the same operation" to Old(-) (or any other starting distribution) is *to use the same magnification factor* β' (and thus the same E: not-E Bayes Factor $1+\beta'$).
  It turns out that if we understand "applying the same operation" in that way, then yes doing the operations in reverse order gives us the same final result, in other words **the operations commute**. This holds in general, even if the operations used different partitions, so long as we identify an operation with the partition it uses, the X: Y Bayes Factors of the operation, with X and Y being each pair of cells of that partition, and we assume the operation is "rigid" with respect to those cells.

- Another idea for "applying the same operation" to Old(-) is that we should *end at the same result* — or at least at the same final values for E and not-E. That is, E should go from Old(E) to New´(E). That will involve a larger change to the value of E than when it merely went from New(E) to New´(E). Moreover, consider what "applying the same operation," understood in this way, would involve when starting from each of these distributions:

  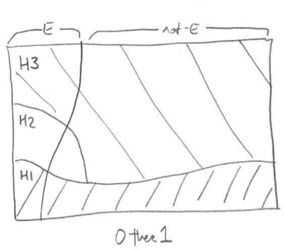

  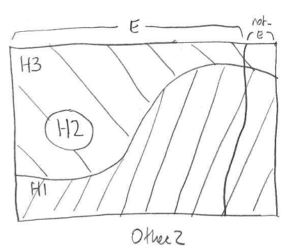

  We'd be applying a *huge* increase to the value of E when operating on Other1(-); and a *decrease* to the value of E when operating on Other2(-). It's not *absurd* to view this "doing the same thing" — since E is always ending at the same value — but neither is it absurd to view it as "doing something different."

  If we understand "applying the same operation " in this way — that is, preserving the final values of the cells, so that New´´(E) = New'(E) and New´´(not-E) = New´(not-E) — then clearly we won't get commutativity. In fact New´´(-) would be the same as New´(-) for all values, and whereas in the first place we went: Old(-) $\rightarrow$ New(-) $\rightarrow$ New´(-), now we'd instead be going Old(-) $\rightarrow$ New´(-) $\rightarrow$ New(-). So our ending distribution would be different.

- Perhaps there are other natural ideas for how to understand "performing the same operation" as when we moved from New(-) to New´(-), but starting with a different distribution. But it turns out that our first choice, as well as being intuitively natural, is on modest assumptions not only sufficient but also necessary for commutativity. (Keeping fixed our assumption that "the same operation" involves using the same partition and being "rigid" with respect to its cells.)

Undermining Evidence

Species of defeating evidence:
- Opposing/rebutting/overturning evidence
- Undermining / undercutting evidence

Many real cases are mixed, but the notion of "pure" undermining evidence seems intelligible
Undermining comes in degrees (even if certain that U, may not entirely neutralize your original evidence E's support of H)
Undermining evidence can itself be defeated/strengthened

Undermining seems to be *ubiquitous* (strategies for generating)
"Confirmational holism" can be understood as:
(i) *All* epistemic justification depends (at least in part) on independent/antecedent justification to believe other things
(ii) *Every belief* can be defeated, or more specifically undermined [**Weak Holism**]
Weisberg: "in general, a belief's empirical justification *is sensitive to* background belief."
Our discussion is just about (ii), and just for empirical beliefs, and even so understood, (ii) is much stronger than needed to generate a problem.

One example of Undermining:
Experiential impact ===> hypothesis E goes up ---> by virtue of bridge premises linking E with H ---> H also goes up
The undermining evidence U threatens the bridge premises.
"By itself" U doesn't speak against H, but when your confidence in H is partly based on support from E via those bridge premises, gaining evidence for U will make H go down.
Call this "downstream" undermining (Pryor 2013: "quotidian undermining")

But what about the support you acquired for E itself: the alleged first impact of your experience. Can *that* support be undermined? Or is undermining always restricted to happening downstream from the proposition E you update on?
We'll rely on Weak Holism only to support the thought that *it's at least possible* for this support for E itself to be undermined.

We get a problem if we try to capture the idea that this can *ever* happen with Jeffrey Conditionalization's representations of your updating.

Three Arguments to Discuss

1. Garber's Complaint
   Given Field's proposal, "…after *nine* repetitions of the *same* rather uninformative experience [seeing a ball in dim light, initially taking you from .3 to .4 for "this ball is blue"], S will become *virtually certain* [credence > .95] that the ball is blue" (Garber 1980, p. 144).

   The underlying issue is that the experiences aren't sufficiently independent to justify how far and how fast the hypothesis "this ball is blue" rises.

   Compare a case where you read a newspaper reporting that H. Reading another copy of the same newspaper? Not evidentially useless (the first could have been a single prank copy your friends planted), but its evidential impact will be negligible. Reading a different paper? More useful, but still not entirely independent. If one paper turns out to be misleading, it's not inevitable but it is likely that others will too, in similar ways.

   When sources of evidence are independent: p(Source2 favors H | H & Source1 favors H) = p(Source2 favors H | H)

   Seeing the same report in multiple copies of the same newspaper, or repeating Garber's casual glance, are extremely far from being independent in this sense.

   This kind of **redundancy** (lack of independence) in one's evidence *operates like a kind of undermining*.

   We'll see that Jeffrey Conditionalization has difficulties capturing undermining effects (when they're operating at the source of one's learning, rather than downstream from it). Garber's complaint is just another manifestation of this difficulty.

2. Christensen's Argument (his 1992, pp. 532-4; also Weisberg 2009, p. 11; also Pryor 2013, p. 127)
   The posterior for E should depend on more than just experiential input, and your priors in E. Your prior credence in U should matter too. But Field's proposal is naturally understood as making the posterior for E depend on just those first factors.

3. Weisberg's Main Complaint (his pp. 14-16; also Pryor 2013, pp. 123-5)
   Suppose U is a pure underminer of the support you will get for E: thus $Old(E|U) = Old(E)$.
   When you learn E, should be that: $Old(E) \leq New(E|U) < New(E)$.

   But Theorem "Rigidity preserves Independence": an update rigid with respect to E cannot take propositions E, U that were independent and make them probabilistically correlated / no longer independent.

References

David Christensen, "Confirmational Holism and Bayesian Epistemology," *Philosophy of Science* 59 (1992)

Marc Lange, "Is Jeffrey Conditionalization Defective By Virtue of Being Non-Commutative? Remarks on the Sameness of Sensory Experiences" *Synthese* 123 (2000)

James Pryor, "Problems for Credulism," in Chris Tucker, ed. *Seemings and Justification: New Essays in Dogmatism and Perceptual Justification* (OUP, 2013)

Jonathan Weisberg, "Updating, Undermining, and Independence," *British Journal for the Philosophy of Science* 66 (2015)