Joyce - How probabilities reflect evidence

Brian T. Miller

October 4, 2017

The most prominent account of the rationality of partial belief states is Bayesianism.

Bayesianism can be regarded as part of a theory of *practical* decision making:

- What should I do?
  - Maximize expected utility
- How do I do that?
  - For each course of action A and each possible outcome of pursuing that course of action, calculate:
    $EV_A = \sum_i P(outcome_i \mid I\ do\ A) \cdot Value(outcome_i)$
  - This presupposes probabilistic beliefs (the $P(outcome_i)$ part)
  - But if my probabilistic beliefs are incoherent then I'll end up doing obviously irrational things (see: Dutch book arguments)
  - Bayesianism provides an appealing theory of coherence[1]

[1] Actually that's complicated. Bayesianism is both a theory of when a set of beliefs is coherent at a moment in time, and it's also a theory of the appropriate response to new propositional evidence, which is a phenomenon that plays out over time; it's a theory of both *synchronic* coherence and *diachronic* coherence. Dutch books are most commonly used to motivate the synchronic bit of the theory, though there are also diachronic versions of the argument.

But Bayesianism can also be regarded as an epistemology. Joyce identifies three components of Bayesian *epistemology*:

Evidential Probability At any time t, a rational believer's opinions can be faithfully modeled by a family of probability functions Ct, hereafter called her credal state, the members of which accurately reflect her total evidence at t.

Learning as Bayesian Updating Learning experiences can be modeled as shifts from one credal state to another that proceed in accordance with Bayes's Rule.

Confirmational Relativity A wide range of questions about evidential relationships can be answered on the basis of information about structural features credal states.

This paper is focused on Evidential Probability.

Probabilities reflect total evidence along three distinct (but related) dimensions:

Balance how decisively the data tells in favor of the proposition... what individual probability values reflect

Weight the gross amount of relevant data available... reflected in the concentration and stability of probabilities in the face of changing information

Specificity a matter of the degree to which the data discriminates the truth of the proposition from that of alternatives [reflecting] the spread of probability values across a credal state (154)

Let's take those in turn.

Balance vs Weight

Here is Keynes:

> As the relevant evidence [for a hypothesis] at our disposal increases, the magnitude of [its] probability may either decrease or increase, according as the new knowledge strengthens the unfavorable or favorable evidence; but something seems to have increased in either case?we have a more substantial basis on which to rest our conclusion...New evidence will sometimes decrease the probability of [the hypothesis] but will always increase its 'weight'. (1921, p. 77)

The intuition here is that any body of evidence has both a kind of valence and a size. Its valence is a matter of which way, and how decisively, the relevant data 'points.' A body of evidence will often be composed of items of data with different valances that need to be compared. It is this 'balance of the evidence' that credences reflect. The size or 'weight' of the evidence has to do with how much relevant information the data contains, irrespective of which way it points. As Keynes emphasized, we should not expect the weight of a body of evidence to be reflected in individual credence values. From the fact two hypotheses have the same credence we can infer that the balance of the evidence for each is the same, but we cannot infer anything at all about the relative weights of the evidence in their favor. (158-9)

Example:

> Four Urns: Jacob and Emily both start out knowing that the urn U was randomly chosen from a set of four urns urn_0, urn_1, urn_2, urn_3 where urn_i contains three balls, i of which are blue and 3 - i of which are green. Since the choice of U was random both subjects assign equal credence to the four hypotheses about its contents: $c(U = urn_i) = 1/4$. Moreover, both treat these hypotheses as statements about the *objective chance* of drawing a blue ball from U, so that knowledge of $U = urn_i$ 'screen offs' any sampling data in the sense that $c(B_{next} \mid E \,\&\, U = urn_i) = c(B_{next} \mid U = urn_i)$, where B_{next} says that the next ball drawn from the urn will be blue and E is a proposition that describes any prior series of random draws with replacement from U. Finally, Jacob and Emily regard random drawing with replacement as an *exchangeable* process, so that any series of draws that produces m blue balls and n green balls is as likely as any other such series, irrespective of order. Use BmGn to denote the generic event in which m blue balls and n green balls are drawn at random and with replacement form U. Against this backdrop of shared evidence, suppose Jacob sees five balls drawn at random and with replacement from U and observes that all are blue, so his evidence is B5G0. Emily, who sees Jacob's evidence, looks at fifteen additional draws of which twelve come up blue, so her evidence is B17G3. What should Emily and Jacob think about B_{next}? (159)

So what will Jacob and Emily's respective beliefs be regarding the probabilities of each urn?

Jacob:

$c(U = urn_0 \mid B5G0) = 0$
$c(U = urn_1 \mid B5G0) = 0.0036$
$c(U = urn_2 \mid B5G0) = 0.1159$
$c(U = urn_3 \mid B5G0) = 0.8804$

Emily:

$c(U = urn_0 \mid B17G3) = 0$
$c(U = urn_1 \mid B17G3) = 0.00006$
$c(U = urn_2 \mid B17G3) = 0.99994$
$c(U = urn_3 \mid B17G3) = 0$

And what should their confidence be that the next ball drawn will be blue? I.e., what's $c(B_{next})$?

For each it should be:

$$\begin{aligned} c(B_{next}) &= \textstyle\sum_i c(U = urn_i) \cdot c(B_{next} \mid U = urn_i) \\ &= \textstyle\sum_i c(U = urn_i) \cdot i/3 \end{aligned}$$

For Jacob that means: $c(B_{next} \mid B5G0) = 0.959$

For Emily that means: $c(B_{next} \mid B17G3) = 0.666626$

Observations:

1. In one sense, Jacob's evidence is better: it speaks more strongly in favor of B_{next}.
   - i.e., the *balance* of Jacob's evidence is further tilted towards B_{next}, in the sense that his credence in B_{next} conditional on *his* evidence is much higher than Emily's credence in B_{next} conditional on *her* evidence
2. in another sense Emily's evidence is better: she has more of if
3. i.e., Emily's evidence has more *weight*

Question: the *balance* of one's evidence for B_{next} shows up in one's credences in the most obvious way: it's just $c(B_{next} \mid E)$. Where does the *weight* show up in your credences?

> ...the weight of the evidence for a proposition X often manifests itself not in X's unconditional credence, but in the **resilience** of this credence conditional on various potential data sequences. A person's credence for X is resilient with respect to datum E to the extent that her credence for X given E remains close to her unconditional credence for X. Note that resilience is defined relative to a specific item of data: a person's belief about X may be resilient relative to one kind of data, but unstable with respect to another. That said, it is usually the case that the greater volume of data a person has for a hypothesis the more resilient her credence tends to be across a wide range of additional data. Our example illustrates this nicely. Even though Jacob's evidence points more definitively toward a blue ball on the next draw, his credence is less resilient than Emily's with respect to almost every potential data sequence, the sole exceptions being those sequences in which *only* blue balls are drawn. In this regard Emily's evidence is better than Jacob's: even though she is not so sure as he is that a blue ball will be drawn, her level of confidence is better informed that his, and so is less susceptible to change in the face of new data. (161)

Specificity

> ...data is less than fully specific with respect to X when it is either incomplete in the sense that it fails to discriminate X from incompatible alternatives, or when it is ambiguous in the sense of being subject to different readings that alter its evidential significance for X. Both incompleteness and ambiguity are defined relative to a given hypothesis, and both are matters of degree. When you are told that Ed is either a professional basketball player or a professional jockey you are given very specific information about the hypothesis that he is an athlete, but somewhat less specific information about the hypothesis that he is a

> jockey. Likewise, if you draw a ball at random from an urn and examine it under yellow light that makes it hard to distinguish blue from green, then finding that the ball looks blue gives you specific information about how it appears in yellow light, but the data is ambiguous with respect to the hypothesis that the ball is actually blue. (171)

Unspecific evidence is prima facie problematic for Bayesianism in cases where the evidence supports two hypotheses equally, e.g. the observation of the greenness of the emerald is consistent with both *the emerald is green* and *the emerald is grue*, which are themselves jointly inconsistent. So what should my posterior credence be in those propositions?

Possible response: equal confidence

This response is often motivated by appeal to the Principle of Sufficient Reason:

PSR: hypotheses... for which there is symmetrical evidence should always be assigned equal probabilities

Problem for PSR: there's a cube with side length between 1cm and 2cm, and my evidence is perfectly symmetrical. What should I believe about the side length?

Possible response: fancy math.

Joyce: that's all nonsense. A uniform distribution of credences over possibilities (as required by PSR) is only rationally obligatory when there's evidence supporting a uniform distribution over possibilities. In cases where PSR might be thought appropriate, no such data is had. So don't follow PSR.

So what should you do?

> the proper response to symmetrically ambiguous or incomplete evidence is not to assign probabilities symmetrically, but to refrain from assigning precise probabilities at all. Indefiniteness in the evidence is reflected not in the values of any single credence function, but in the spread of values across the family of all credence functions that the evidence does not exclude. This is why modern Bayesians represent credal states using sets of credence functions. It is not just that sharp degrees of belief are psychologically unrealistic (though they are). Imprecise credences have a clear epistemological motivation: they are the proper response to unspecific evidence. (171)

So how is the *specificity* of evidence reflected in your credences?

> The specificity of the evidence is reflected in the spread of credence values for the proposition across the believer's credal state. (176)