HPSC 1001/1901/2101/2901

WHAT IS THIS THING CALLED SCIENCE?

Semester 2, 2020

Lecture 20: Bayesianism and Modern Theories of Evidence - 2

The combination of ideas we are looking at:
1. Degrees of belief,
2. Interpreted as probabilities,
3. Updated via Bayes' theorem.

Yesterday: 1, 2, and part of 3. Finish 3 today and then look at where we end up.

... 3. Updated via Bayes' theorem

Introduce *conditional probability*. The probability of one thing *conditional on* another thing – or *assuming* another. P(X | Y): probability of X conditional on Y.

In a conditional probability P(X|Y), we *ignore* whether you are certain or uncertain about Y. Just *assume* Y, and ask about the probability of X given Y.

P(X|Y) is not in general the same as P(Y|X). Though the numbers might match in a particular case.

Bayes' theorem. An 18th century result. Controversial for many years. Always looked like it *might* help with evidence. But seemed to also lead to problems. It gets a new role within "subjectivist" interpretations of probability. Hence: "subjectivist Bayesianism."

An example. Suppose you think:

P(Pat is at the party) = 1/2

P(Pat's car is outside | Pat is at the party) = 0.8

P(Pat's car is outside | Pat is not at the party) = 0.01

P(Pat is at the party | Pat's car is outside) = ?

Bayes' theorem:

Applied to evidence:

How likely are we to see e, if h is true?

$$P(\text{hypothesis } h|\text{ evidence } e) = \frac{P(h) * P(e|h)}{P(e)}$$

Without considering the new evidence, what is the overall or "prior" probability of h?

How likely are we to see e, overall? (Whether or not h is true). This will get explained further in a moment.

More on P(e): We want to take into account how likely e is if h is true, and also how likely e is if h is false.

$$P(e) = P(e|h) * P(h) + P(e|\text{ not } h) * P(\text{not } h)$$

What is P(h)? In older views of probability, there was much discussion of whether these "prior probabilities" of theories make sense at all. What do they measure? Maybe Bayes' theorem cannot be used in philosophy of science, because those prior probabilities of theories make no sense.

If you are a subjectivist, following Ramsey, this is no big deal. A "prior probability in h" is just your degree of belief in h before the evidence comes in. Fine.

Now we have all we need to solve some problems of evidence.

Hypothesis h is that Pat is at the party.
The evidence e is seeing Pat's car.

Do the maths from the case above:

$P(h|e) = (0.8)(.5)/((0.8)(0.5) + (0.01)*(0.5)) = \text{(approx) } 0.99$

When I did the example I did not think the number would come out quite so high. That 0.01 really has an effect; it would be very hard for her car to be out there without her being inside.

The evidence *confirms* the hypothesis: $P(h|e) > P(h)$

$P(h|e)$ is called the *posterior* probability of h.
$P(h)$ is called the *prior* probability of h.
$P(e|h)$ and $P(e|\text{not } h)$ are both called *likelihoods*.

Suppose you thought she was very unlikely to show up:
P(Pat is at the party) = 0.1.
Then:
P(Pat is at the party | Pat's car is outside) =(approx) 0.89

Still very high! Again, that 0.01 really has an effect.

Two other parts of this model of belief and evidence:
How change looks over many episodes.

Today: you have your P(h). Degree of belief in h.
You also have views about what you are likely to see if h is true, and what you are likely to see if h is false.
P(e|h), P(e| not h), etc....
You can calculate P(h|e), even before you see e.
Then you see e.

You *update* your degrees of belief in the light of this new evidence.
$P(h|e)$ becomes your *new degree of belief* in h.
It becomes your new $P(h)$.
Then you can do it all again.

We can write it out more exactly by using subscripts to indicate which days we are talking about:

Assume we see e at the end of today:

$$P_{tomorrow}(h) = P_{today}(h|e)$$

More fully:

$P_{tomorrow}(h) = P_{today}(h|e) = P_{today}(h)* P_{today}(e|h) / P_{today}(e)$

Today's posterior is tomorrow's prior.

Confirmation: evidence e *confirms* h if and only if $P(h|e)>P(h)$.

The overall picture: you start out with freely chosen degrees of belief in all sorts of possible theories.

The dinosaur case: maybe early mammals ate the dinosaur eggs? Maybe there was a slow process of climate change? Maybe an asteroid hit the Earth?
You might give these roughly equal initial degrees of belief (or you might not).

As evidence comes in, you update those initial degrees of belief. That part of the story is not so "free". Given your initial degrees of belief, and your evidence, there is only one move you should make next.
In the dinosaur case: if you see lots of iridium at 65 mya in the rock layers, you increase your degree of belief in the

Alvarez hypothesis about an asteroid hitting the Earth. Other competing hypotheses must lose some of your confidence.

* From a Bayesian point of view, the differences between evidence in science and everyday life can be minor differences (as the old empiricists thought).
Hypothesis h can be a very complicated scientific theory, or it can be the hypothesis that your friend is at a party.
Evidence e could be data from a supercollider experiment, or iridium in the rock layers, or an observation of a car outside a house.

How this relates to the "science and values" discussion. There I said: you should work out your degrees of belief, using the evidence, and then use the degrees of belief along with your preferences (or goals) to work out what to do. Now we have a theory of the degree of belief side.

Should I take my umbrella? Should I vaccinate all children? Should I stop mining and burning coal? All these cases can be handled in (roughly) the same way.

Suppose you have a range of possible actions, A_1, A_2, etc.

Expected value of action $A_1 = P(h_1)V(A_1|h_1) + P(h_2)V(A_1|h_2)$

And so on, for action A_2 and any others.
Choose the act that gives you the highest expected value.

"$V(A_1|h_1)$" looks a bit like a conditional probability, but it's not; it's a value or payoff – the value of performing A_1 in a situation where h_2 is true (eg., carrying an umbrella all day when it doesn't rain).

Assessing Subjectivist Bayesianism
What *kind* of theory is this?
A theory of rational belief changes by an individual.

What counts as 'rational'? Following the rules of probability. Why is that a good idea? Ramsey argued that if you don't, you will get into trouble with your actions. But only in very special circumstances.

The overall picture: you start with some initial P(h), and update as evidence comes in. Where you end up depends on where you start. A 'good' or 'justified' degree of belief is very much dependent on where you start.

This is a very different way of talking about evidence.

Forget the idea that:

"*The* probability that Darwinism is true is 0.9"

Or even: "Given our evidence, the probability that

Darwinism is true is 0.9"

Instead: "A rational person who starts out here... and sees this evidence.... will end up with a degree of belief in Darwinism of 0.9."

The prior probabilities at the start of the process are a "free choice." So where you end up depends on where you start. If you start somewhere crazy, you will end up somewhere crazy? Not always.

"Washing out" results -- see how little the differences in prior probabilities did in the Pat's car case above. The two people ended up in a similar place.
(But see T&R p. 209 for some further comments about this.)

I said this is a "theory" of belief change. But it seems to have some weird features in the set-up. You are supposed to start out with all the hypotheses you will ever consider listed and given some initial degree of belief. (Analogy in tutorial yesterday -- it is like a stock market situation where every company you might invest in is listed at the start, and none can be added.)

Well - it need not be quite like that. You might have in your initial list of hypotheses a hypothesis like this:
"Something I haven't thought of yet."
But: if you include something like this, it makes it harder to update degrees of belief using Bayes' theorem.
Because to use the theorem, you need to know P(e).
That now requires knowing: P(e|Something I haven't thought of yet).

Why is that? See above for our $P(e)$ in the party case:

$$P(e) = P(e|h) * P(h) + P(e|\text{ not } h) * P(\text{not } h)$$

That is a simple case. Often, you will have several hypotheses to consider:

$$P(e) = P(e|h_1)*P(h_1) + P(e|h_2)*P(h_2) + P(e|h_3)*P(h_3) \ldots$$

And now we are saying that one of these hypotheses is "Something I haven't thought of yet."

* How we might think of the whole situation:
Subjectivist Bayesianism is a very good *model* of rational belief change.
A model: it simplifies some things, but captures the basics very well.
Those basics: degree of belief, its relation to evidence, and how belief relates to action.