

Some Old Objections to Howson and Urbach's Defence of Bayesianism Revisited

JOHN WORRALL

Colin memories

So all three of us were of course colleagues in the period that Colin and Peter were writing their marvellous book.

Of course, admiring the book did not mean that I had no objections and in fact I had two main issues with their view back then, neither of them, I'm afraid, the least bit original: on the contrary, they could be termed the "bog standard" objections.

First, I was one of those annoying people who couldn't help thinking that classical statistics could not be as fundamentally and multiply flawed as Colin and Peter argued it was, and yet at the same time be so widely accepted and applied in serious disciplines like clinical medicine and psychology.

This was always a sociological rather than a fully thought-through logical objection of course. And on this issue, having studied their arguments more carefully, I have for many years now been a complete convert. Classical statistics *is* just as fundamentally and multiply broken-backed as they consider it to be – indeed, as, for any fair-minded observer, they *demonstrate* it to be. Much of the trouble stems from the wrongheaded decision to draw an arbitrary line across a continuous spectrum and 'pretend' that a probabilistic theory has been falsified when an event occurs that is deemed sufficiently improbable on the assumption that that theory is true. That way lies the idiocy of thinking there is such a thing as the "lottery paradox" and lots of other unnecessary troubles besides: "if you are modelling some situation as probabilistic, then stick with probabilities" should surely be the rule.

My *second* erstwhile objection is however a different matter; and indeed is still current. This is that (and if you must yawn, and no one would blame you, could you please yawn discreetly?): subjective Bayesianism, unless supplemented by some additional principles, is not strong enough to provide a satisfactory account of the objectivity or rationality of science. This problem is usually attributed to the fact that what we might call "pure" or "basic" Bayesianism places no constraints on prior probabilities; but, as we shall shortly see, I think there is more to the problem than that.

Being such a standard objection, Colin and Peter, naturally have responded to it. A central plank of their response is that not every view about the merits of rival scientific theories in light of the evidence available at a given time is objectively correct or incorrect: there is room for reasonable disagreement and it is a virtue of the Bayesian approach, rather than a vice, that it provides that room by allowing for an element of subjectivity via the priors.

Taken across science as a whole and looking at any stage of development - particularly its early stages, this is undoubtedly correct. There was, for example, no single correct view as to which of the corpuscular and wave theories of light was better supported by the evidence available in the 1660s when Newton and Hooke were debating the issue. Or rather, since Hooke's line was always a sinker and Newton's not a whole lot better, the objectively dictated view was that neither theory had any real empirical support and hence, so far as those two great scientists had different theory-preferences, those preferences reflected prior, subjective elements.

But the dynamic duo go way too far, alleging

“[S]ome subjective element exists in *all* scientific appraisal and it is a merit of this [Bayesian] theory that where it occurs it is signalled explicitly, not concealed from view” (3rd ed, p.265; my italics)

Really: in *all* scientific appraisals? Can it really be reasonable, for example, in view of all the optical evidence now available, to prefer Newton's corpuscular theory of light to current theories based on quantum electrodynamics? Admittedly, there are (probably) no real people whose degrees of belief would reflect that preference. But there are, for example, real people - many of them - , who at least call themselves scientists and whose degrees of belief reflect a preference for young earth creationism over theories based on cosmology and natural selection as an account of the current flora and fauna of the earth. Can those beliefs be even possibly rational or scientific?

Now you might think that young earth creationism, with its claim that the earth and indeed the whole universe was created in 4004BC or so, is simply empirically refuted, since we know, i.e. we have unassailable evidence, that there are countless things in the universe orders of magnitude more than 6000 or so years old.

Hence young earth creationism has a Bayesian probability of zero and hence represents no problem for the view.

But enter my favourite *ad hoc* theorist. This is Philip Henry Gosse and this is a copy of his masterpiece *Omphalos*. Gosse was fascinated by the question of why, given

the unorthodox nature of his birth-process, Adam, whom Gosse took to be a real historical figure, should have had a navel.

(Incidentally it is not clear exactly on what basis Gosse believed that Adam did have a navel. *Genesis* is my favourite bedtime reading and I can assure you that there is no mention of a navel there. (It does famously say that God “created man in his own image; in the image of God created he him”. But surely the explanation can’t be that God has a navel - that would be deeply weird and even if someone did endorse this, along with God’s decision to create man in His own image, as the explanation of Adam’s outer abdominal structure - talk about wrapping a mystery in an enigma!)

Maybe Gosse just thought that there are so many painters in the history of art who painted Adam with a navel that they couldn’t all have got it wrong.)

In any event, Gosse was fascinated by the problem and, after producing quite an amount of somewhat turgid prose in response to it, he comes to the conclusion (essentially): “well, why *shouldn’t* God have decided to give Adam a navel?”. After all (implicitly), He moves in mysterious ways. And Gosse promptly uses this same idea as an explanation for why many things in the universe appear to be much older than seems compatible with the Creation story/theory : God created the universe in roughly 4004BC but He created many things looking as if they were already very old, created the rocks of the earth with what look like marks made by the skeletons of extremely ancient and now extinct species, but are in fact just random marks that happened to please Him, and so on.

So creationism is not out of the game as having probability zero. The much-rehearsed problem then, of course, is that, supposing no constraints at all on prior probabilities, an agent may have started with an overwhelmingly high prior for young earth creationism, have Bayesian-updated away in approved fashion on the incoming evidence and yet still place a substantial (posterior) probability on the creationist theory, one even perhaps higher than the probability s/he assigns to rival Darwin-based views; and that just can’t be right. No account of the basis for scientific reasoning can, it seems to me, be correct if it allows this even as a possibility.

Nor can the various lauded limiting-case results on their own help: even if it can be shown that a Bayesian creationist is destined, in the indefinite long run, to see her degrees of belief become aligned with others who currently hold more orthodox views, this is not enough: a Bayesian who ascribes a substantial probability to young earth creationism is wrong, objectively wrong, NOW.

Nor is the problem of someone’s starting out by assigning what to any sensible observer would seem a ridiculously high prior to a theory, the only problem.

It seems to me that there are *two other* sources of subjectivity in the “pure” subjective Bayesian view – though it does worry me that, while they seem so obvious, they have been given no attention: of course, I am not fully up-to-date so there may have been attention that I missed.

The *first* is that the pure view places no objective constraints on what counts as evidence. Evidence for an agent is, it seems, whatever s/he comes to assign a probability of one – exogenously, independently of probabilistic calculation. Now in any remotely sensible Bayesian reconstruction of some episode in science, the evidence will be taken to be things like Newton’s observations of the spectrum produced by passing sunlight through a prism, the outcome of the two-slit experiment with monochromatic light, or the outcome of the Michelson-Morley experiment or whatever. But the requirement that *only* such things count as evidence would be an added assumption of the kind that Colin in particular eschewed (for reasons we will come to). Without such an assumption, however, subjective Bayesianism is even more clearly at odds with what seem to me compelling normative judgments about scientific rationality.

To illustrate: I once made the mistake of agreeing to help out Adolf Grünbaum who had contracted to write a chapter on Science and Religion for a volume on the Philosophy of Religion, but found himself too busy to write it. The article I eventually produced led to some fairly furious emailing with philosophers inspired by Alvin Plantinga, Peter van Inwagen and the like who call themselves ‘new epistemologists’. They told me that for them, the existence of God was a straightforward evidential matter: they see God in, for example, every rosebud. But surely an agent cannot count as scientific/rational if they regard ‘I see God in this rosebud’ (whatever that is supposed to mean) as having probability one and therefore as evidential. So it seems clear that we at least need to add constraints on what can count as evidence to the basic Bayesian view if we are to approach a defensible account of scientific/rational reasoning.

So we have the subjectivity of the priors, and the subjectivity of evidence, and there is, it seems to me, *a third* source of over-subjectivity concerning background knowledge.

Taking background knowledge into account is, of course, essential to sensible confirmation judgements. For example, as Colin and Peter point out:

“To talk about e being evidence relevant to h obviously requires a background of fact and information against which e is judged to be evidence. A large dictionary found in the street is not in itself evidence either for or against the hypothesis that Smith killed Jones. Relative to the background information that Jones was killed

with a heavy object, that the dictionary belonged to Smith, and that blood found on the dictionary matches Jones's, it is." (3rd ed, p.298)

Or again, nowadays - given that background knowledge includes observations of black swans - no number of observations of white swans can supply any confirmation of the hypothesis h that "all swans are white". For any e , $p(h/e \& B) = p(h/B) = 0$, so no confirmation. For pre-Cook Europeans with their different background knowledge, observation of individual white swans *could* confirm "all swans are white".

And indeed for a fully satisfactory analysis of scientific reasoning, background knowledge must be taken as containing already massively confirmed generalisations and theories as well as lots of singular pieces of evidence.

But any adequate account of scientific reasoning is surely going to have to impose constraints to ensure that what gets into background knowledge *and* what counts as a justified change in background knowledge are evidence-based affairs.

"Pure" subjective Bayesianism, however, - requiring only that at each stage, an epistemic agent's degrees of belief satisfy the probability calculus and that thereafter she updates on the accumulating evidence - explicitly spurns any such constraints. Instead, so far as I can tell, a Bayesian creationist who starts with an extreme prior for creationism but then updates correctly leading to a rapidly diminishing posterior can, consistently with the subjectivist view, nonetheless find at any stage that her background knowledge has changed (perhaps via some nocturnal interchange with God, though there is no requirement that she explain the change) and that that change involves a new round of bets involving a new set of priors - priors that perhaps takes her back to a very high probability for creationism. Such a person surely cannot be considered to be assigning her degrees of belief in a rational, scientific way.

So three sources of subjectivism point to the need for extra constraints if Bayesianism is to provide a fully adequate account of scientific reasoning. Or so it seems to me. For Colin, on the contrary, (a) no more objectivity is required than that supplied by basic Bayesianism¹ and (b) even if you thought that more objectivity is required, you couldn't have it. On (a) I believe he was just wrong for the reasons already given; (b) is interesting.

In a number of articles and in the last chapter of the 3rd edition of the book, Colin argued *first* that the basic Bayesian constraints are dictated by logic alone: logic can

¹ And I agree with him that any suggestion (does anyone really make it?) that pure Bayesianism supplies *no* objective constraints at all would be substantially mistaken.

be thought of as the study of consistency, and the subjective Bayesian requirement that degrees of belief be probabilistically consistent is completely on a par with the requirement of deductive consistency between sentences. And *secondly* that attempts to develop a fully objective Bayesianism – attempts, it should be noted, not just to put constraints on acceptable priors but to specify *uniquely correct* priors – all run into sand: the most famous such attempt being based of course on the principle of indifference - a principle that, no matter how much plausibility it initially might appear to have, on analysis, not only cannot be defended as in some sense ‘logically’ required, but also turns out to be formally inconsistent.

Well, we have with us *the* objective Bayesian (Jon Williamson) who won’t be happy, but speaking personally I *am* happy to accept that any attempt to produce an objective Bayesianism that (a) specifies uniquely correct priors and (b) is based only on content-neutral, structural considerations has failed and no other contender is on the horizon. It doesn’t follow, however, that we can’t defend some weaker constraints of a non-logical kind that deal with what I have argued is the clearly unacceptable amount of subjectivity licensed by the subjective Bayesian account.

Discussing an issue close to my heart: that of whether we can justify a preference perhaps on grounds of simplicity for a theory with fewer adjustable parameters over one with more adjustable parameters (of course *ceteris paribus*), Colin comes to the conclusion that really the issue is simply that of which theories seem plausible. “Ultimately, the criterion counting above all is plausibility in the light of background knowledge; to the extent that simplicity is a criterion, it is to the extent that it supervenes on prior probability, not the other way round.” (292)

Let’s suppose that this is correct, still the issue seems obvious: for instance, an agent who consistently found a theory more plausible the more adjustable parameters it contained will count, on subjective Bayesianism, as just as rational as the orthodox scientific thinker who always looks for fewer adjustable parameters. Surely in order to develop an adequate account of scientific reasoning, the Bayesian needs to add an account of what can reasonably, scientifically be taken to be plausible in the light of background knowledge rather than just leave it completely to subjective taste?

As noted, Colin stoutly resisted this suggestion. Why? Well, I couldn’t find anything in print, but from conversations I had with him and also from an interchange he had with Clark Glymour after a talk at a conference in Warsaw many years ago, I think Colin’s position here was basically as follows.

Any further constraints beyond basic Bayesianism would rely on substantive assumptions. Only deductive logic plus basic Bayesianism count as logical - attempts to develop an objective Bayesianism based on extra quasi-logical or

structural elements all failing. Hence further constraints would rely on substantive assumptions, but then what happens if someone challenges those assumptions? It is difficult to see how to mount a rational defence of principles that are themselves alleged to (in part) constitute rationality; and even if such a defence could be mounted, it would clearly itself have to rely on further substantive assumptions and what if *those* were challenged?

That is, Colin was motivated here by the dreaded *infinite regress problem*. Or, as I see it, since the only way to halt a potential infinite regress is to dig in your heels and be dogmatic: by the *problem of necessary dogmatism*. Suppose, for example a Bayesian developed some constraint which, while certainly not specifying unique priors nonetheless somehow ruled out some prior distributions as overly prejudicial, so that an agent who starts by assigning a prior of 0.9999999999999999 for creationism is declared irrational. And suppose someone asks why she shouldn't be allowed such initial priors and the answer just is 'rationally you just can't – it's obvious'. This doesn't at all sound like a comfortable position for a rationalist. A rationalism based on dogma - what sort of rationalism is that?!

So, Colin's view was that deductive logic plus basic Bayesianism is logic, which (implicitly) cannot be questioned; anything further would hit the infinite regress/necessary dogmatism problem and hence could not be rationally defended; so rationality is restricted to deductive logic plus basic Bayesianism and if that sanctions too much subjectivism for your taste, then get used to it and adjust your tastes: no restrictions on that subjectivism can be rationally justified.

This might sound persuasive but I am not convinced. For one thing, logic can be, and has been questioned. Indeed the whole thrust of Lewis Carroll's "What the Tortoise said to Achilles" is precisely that modus ponens, for example, can be coherently questioned and that if It is questioned then it can only be justified circularly in a way that assumes modus ponens (of course this assumption will not generally be explicit – but a little digging will reveal that it is implicitly invoked). So either an infinite regress, or you cut it off, right at the beginning by asserting that we just know that Modus Ponens is truth-preserving.

There are some philosophers who, in my opinion, ought to know better – Peter Lipton and Stathis Psillos are two of them - who have tried to draw a dividing line amongst circular arguments between those that are viciously and those that are somehow "virtuously" circular. Virtuously circular arguments can legitimately be taken to be convincing. Some arguments justifying certain sorts of induction are admittedly themselves inductive, but this may somehow be okay: the arguments are circular alright but somehow non-viciously so.

But this is just hot air in my view. Forget any qualifying adjectives: a circular argument is just circular. Trying to develop a notion of a circular argument that is somehow not viciously circular is just a way of avoiding admitting that you are happy to assert certain things - whether deductive principles or certain kinds of 'inductive' ones - dogmatically.

I am in fact perfectly comfortable with asserting that we know that modus ponens is truth-preserving, that massively skewed priors are, in many cases, unacceptable, that predictive empirical successes, *ceteris paribus*, are more confirmatory than *ad hoc* accommodations. There is nothing more basic that you could hope to deduce modus ponens from and maybe the same can be said for appropriate "inductive" principles.

Conclusion: if the Bayesian insists on sticking to the "pure" or "basic" position then she sanctions way too much subjectivity to provide an adequate account of scientific reasoning, and, moreover she only fools herself if she thinks "dogma" is thereby avoided. Sounds like she should be working on developing defensible further assumptions that will take her toward a more defensible account of scientific reasoning via a more (though not fully) objective Bayesianism.