Some final remarks on Williamson's defence of "In Defence of Objective Bayesianism"

Olle Häggström

It has sometimes been said that to be an author of a book is like being the parent of a child. When a stranger complains about your child, this tends to trigger emotion and anger, to the point where it prohibits rational discussion. As a book reviewer, I have sometimes noticed similar reactions from authors. For a prime example, see Johnson's [J04a] response to my book review [H04a] (plus the further correspondence in [H04b] and [J04b]). So it should not have come as a surprise to me that Jon Williamson's response [W11] to my review [H10] of his book "In Defence of Objective Bayesianism" [W10] is distressed, and that, in his eagerness to immediately prove me wrong, he forgets to attempt to understand what I am actually saying. Let me turn immediately to each of the four points that he raises in [W11].

1. Uniform distribution on a large but finite state space

A cornerstone in the epistemology that Williamson [W10] promotes is that our beliefs should adhere whenever possible (i.e., unless constrained by evidence) to uniform distribution. That this is problematic or impossible for infinite state spaces is well known, so in my review I choose instead to point out how bad a choice it can be in cases with a large but finite state space:

> Consider the following image analysis situation. Suppose we have a very fine-grained image with $10^6 \times 10^6$ pixels, each of which can take value black or white. The set of possible images then has $2^{10^{12}}$ elements. Suppose that we assign the same probability $1/2^{10^{12}}$ to each element. This is tantamount to assuming that each pixel, independently of all others, is black or white with probability 1/2 each. Standard probability estimates show that with overwhelming probability, the image will, as far as the naked eye can tell, be uniformly grey. In fact, the conviction of uniform greyness is so strong that even if, say, we split the image in four equally sized quadrants and condition on the event that the first three quadrants are pure black, we are still overwhelmingly convinced that the fourth quadrant will turn out grey. In practice, this can hardly be called unbiased or objective.

Here, I could have written "Standard probability estimates show that *the agent believes that* with overwhelming probability", but chose to omit the four italicized words because I considered it to be clear from the context (and, for instance, from the phrase "the conviction of" in the very next sentence) that the entire discussion is about the agent's belief. I am sure that all readers of my review understood this – except for Williamson [W11], who, in his desperation to find some error to strike down upon, read the omission of those four words as an indication that I hadn't understood the distinction between epistemic and empirical probabilities.

Not that this distinction really matters in this case. My example shows that uniform distribution can be a horribly bad choice of prior distribution – regardless of philosophical subtleties in how to interpret the notion of probability. To invoke the distinction *epistemic vs empirical* against my example is a category error, a bit like trying to understand the inflation rate in the US economy by studying the chemical properties of dollar bills.

2. Language dependence

Moving on to smaller state spaces, I wrote [H10] the following.

> In his first chapter, Williamson describes a situation where a physician needs to judge the probability that a given patient has a given disease S. All the physician knows is that there is scientific evidence that the probability that a patient with the given symptoms actually has disease S is somewhere in the interval $[0.1, 0.4]$. Williamson's suggestion is that the physician should settle for $P(\text{ill}) = 0.4$, because this is as close as he can get to uniform distribution $(0.5, 0.5)$ on the space $\{\text{ill}, \text{healthy}\}$ under the constraint given by the scientific evidence. [...]
>
> Rather than giving the whole list of objections that come to my mind, let me restrict to one of them: what Williamson himself calls *language dependence*. Let us suppose that we refine the crude language which only admits the two possible states "ill" and "healthy" to account for the fact that a healthy person can be either susceptible or immune, so that the state space becomes $\{\text{ill}, \text{susceptible}, \text{immune}\}$, and Williamson's favored estimate goes down from $P(\text{ill}) = 0.4$ to $P(\text{ill}) = 1/3$. By further linguistic refinement (such as distinguishing between "moderately ill", "somewhat more ill", "very ill" and "terminally ill"), we can make $P(\text{ill})$ land anywhere we wish in $[0.1, 0.4]$. How's that for objectivity?
>
> Williamson is aware of the language dependence problem and devotes Section 9.2 of his book to it. His answer is that one's

> language has evolved for usefulness in describing the world, and may therefore itself constitute evidence for what the world is like. "For example, having dozens of words for snow in one's language says something about the environment in which one lives; if one is going to equivocate about the weather tomorrow, it is better to equivocate between the basic states definable in one's own language than in some arbitrary other language" (Williamson, p 156–157). This argument is feeble, akin to noting that all sorts of dreams and prejudices we may have are affected by what the world is like, and suggesting that we can therefore happily and unproblematically plug them into the inference machinery.

Williamson's [W11] response:

> It is sufficient to point out here that the analogy is a false one: languages are not as ephemeral as dreams or prejudices. The language one uses in a particular operating context is rather tightly constrained by the context itself – whether it is the language of a baker or a carpenter or a molecular biologist. For sure, in fictional contexts we can invent gobbledygook terms that are remotely related to reality but that does not apply to day-to-day languages or scientific languages – our terms in these latter languages generate what we consider to be the basic possibilities. Language dependence, then, is not obviously problematic and it is incumbent on anyone who thinks otherwise to come up with realistic cases that demonstrate otherwise.

It is plain incomperhensible to me how Williamson can claim that "language dependence [...] is not obviously problematic" and how he can request that "anyone who thinks otherwise to come up with realistic cases that demonstrate otherwise", in response to the example with the linguistic refinements of the notions "ill" and "healthy". There, right in front of his own eyes, is the case (realistic or not – it is Williamson's own scenario) that "demonstrate[s] otherwise".

In fact, Williamson's entire defence against language dependence can be dismissed on the grounds that it is based on a false identification between language in the precise technical meaning he has given the term (namely the probability space on which the prior is defined) and our everyday meaning of the term. It is just a very vague analogy, and the idea that only some fixed finite number of nuances of "ill" and "healthy" should be expressible in (say) English is simply ludicrous.

3. Dynamic Dutch booking

The most surprising aspect of Williamson's favored variant of Bayesianism is that he rejects the use of Bayesian conditionalization (i.e., transforming the prior into a posterior by conditioning on the available evidence). Since elsewhere in his book he relies heavily on so-called Dutch book arguments, he must respond to the well-known result of Teller [T73] that anyone who deviates from Bayesian conditionalization is susceptible to a sequential Dutch book. Williamson's respose to this is to reject the relevance of sequential Dutch book arguments on the grounds that

> in certain situations one can Dutch book *anyone who changes their degrees of belief at all*, regardless of whether or not they change them by conditionalization. Thus, avoidance of Dutch book is a lousy criterion for deciding on an update rule. [W10, p 85, emphasis in the original]

This, however, is plain false for the case of the Bayesian conditionalizer, as is easily shown by a standard martingale argument. The "certain situations" that Williamson refers to is one that simply does not happen to a Bayesian conditionalizer, namely the following:

> Suppose it is generally known that you will be presented with evidence that does not count against θ, so that your degree of belief in θ will not decrease. [W10, p 85]

In [H10], I point out the impossibility for a Bayesian conditionalizer to end up in this situation, because

> as a Bayesian conditionalizer I would never find myself in a situation where I know beforehand in which direction my update will go, because then I would already have adjusted my belief in that direction.

Williamson's response shows that he doesn't understand: He claims that it *is* possible for the Bayesian conditionalizer to end up in his scenario

> because it is not a foregone conclusion that one's degree of belief will increase in the light of the new evidence – it could stay the same. So there is no reason why my scenario should not be viewed as one in which a conditionalizer might find herself. [W11]

I actually thought my argument in [H10] was simple enough that I would not have to deal explicitly with the case of a possible zero change in belief, but all right, let me be more explicit. Write θ for the event which my (i.e., the Bayesian conditionalizer's) belief in will not decrease in the light of the new evidence, write q for the probability that I assign to θ right now, and write q^* for the probability that I will assign to θ after being presented with the new evidence. Finally, write $P(q^* > q)$ for the probability (according to my own current belief) that my belief in θ will strictly increase when I am exposed to the new evidence. There are now two cases to consider, namely (a) $P(q^* > q) = 0$, and (b) $P(q^* > q) > 0$. In case (a), my belief in θ will not change, and Williamson's [W10, p 86] construction of a sequential Dutch book fails. In case (b), an elementary calculation shows that the expected value $E[q^*]$ (calculated under my current belief) satisfies $E[q^*] > q$. But this situation cannot arise, because as a Bayesian conditionalizer I should already have updated q to coincide with $E[q^*]$. This is simply the martingale property of conditional probabilities. So neither case (a) nor case (b) of Williamson's scenario poses a sequential Dutch book threat to the Bayesian conditionalizer. QED.

4. The Bacchus–Kyburg–Thalos example

As a further argument against Bayesian conditionalization, Williamson invokes the follwing example, originally due to Bacchus, Kyburg and Thalos [BKT90].

> Suppose A is 'Peterson is a Swede', B is 'Peterson is a Norwegian', C is 'Peterson is a Scandinavian', and ε is '80% of all Scandinavians are Swedes'. Initially, the agent sets $P_{\mathcal{E}}(A) = 0.2$, $P_{\mathcal{E}}(B) = 0.8$, $P_{\mathcal{E}}(C) = 1$, $P_{\mathcal{E}}(\varepsilon) = 0.2$ and $P_{\mathcal{E}}(A \wedge \varepsilon) = P_{\mathcal{E}}(B \wedge \varepsilon) = 0.1$. All these degrees of belief satisfy the norms of subjectivism. Updating by [maximum entropy] on learning ε, the agent believes that Peterson is a Swede to degree 0.8, which seems quite right. On the other hand, updating by conditionalization on ε leads to a degree of belief of 0.5 that Peterson is a Swede, which is quite wrong. [W10, p 80]

And here is how I respond to the example:

> Here Williamson obviously thinks the evidence ε constrains the probability of A to be precisely 0.8. This is plain false – unless we redefine C to say something like "Peterson was sent to us via some mechanism that picks a Scandinavian at random according to uniform distribution, and we have absolutely no other information about how he speaks, how he dresses, or anything else

> that may give a clue regarding his nationality". But this is not how the problem was posed.
>
> Suppose however for the sake of the argument that ε does have the consequence that Williamson claims. Then in fact the choice of prior is incoherent, because $P_{\mathcal{E}}(A \wedge \varepsilon) = P_{\mathcal{E}}(B \wedge \varepsilon) = \frac{1}{2}P_{\mathcal{E}}(\varepsilon)$ *means* that given ε, the odds for Peterson being Swedish or Norwegian are fifty-fifty. [H10]

And Williamson responds in return:

> Finally, Haggstrom claims that a specific probability function that I appeal to (in an example due to Bacchus, Kyburg and Thalos) is incoherent, i.e., ill-defined. This is simply not true: it is a well-defined probability function [...]. Haggstrom renders it incoherent by adding further information that was neither present nor required in the original example. [W11]

This response from Williamson is silly beyond belief. As the reader can plainly see, I do *not* claim that the BKT probability function is incoherent. All I claim is that it doesn't have a certain property Q that Williamson claims it has, namely that upon learning evidence ε we should be compelled to assign probability 0.8 to the proposition that Peterson is a Swede. Now, as the reader can also see, Williamson provides no justification whatsoever for for property Q. In my wish to understand how Williamson was thinking, I was forced to speculate as to what implicit assumption might have compelled Williamson to conclude Q, and suggested that it might have been a certain mechanism M concerning how Peterson came to us. But note that it is Williamsons's unwarranted claim Q, rather than my speculation M, that I claimed would render the BKT probability distribution incoherent.

References

[BKT90] Bacchus, F., Kyburg, H. and Thalos, M. (1990) Against conditionalization, *Synthese* **84**, 475–506.

[H04a] Häggström, O. (2004) Ett paradigmskifte i matematiken? *Svenska Matematikersamfundets Medlemsutskick*, May 15.
`http://www.math.chalmers.se/~olleh/skolans_sak/paradigmskifte.html`

[H04b] Häggström, O. (2004) Svar till Claes Johnson.
`http://www.math.chalmers.se/~olleh/skolans_sak/paradigmskifte_2.html`

[H10] Häggström, O. (2010) How objective is objective Bayesianism – and how Bayesian?, *EMS Newsletter*, December issue, pp 59-61.
`http://www.math.chalmers.se/~olleh/Williamson.pdf`

[J04a] Johnson, C. (2004) Johnson granskar Häggström/Svenska matematikersamfundet, *Svenska Matematikersamfundets Medlemsutskick*, October 15. `http://www.math.chalmers.se/~olleh/skolans_sak/Johnson_granskar.pdf`

[J04b] Johnson, C. (2004) Svar till Olle Häggström. `http://www.math.chalmers.se/~olleh/skolans_sak/Johnson_granskar_2.pdf`

[T73] Teller, P. (1973) Conditionalization and observation, *Synthese* **26**, 218–258.

[W11] Williamson, J. (2011) Response to Professor Haggstrom's review of "In defence of objective Bayesianism", *EMS Newsletter*, March issue, pp 58-60. `http://www.ems-ph.org/journals/newsletter/pdf/2011-03-79.pdf`

[W10] Williamson, J. (2010) *In Defence of Objective Bayesianism*, Oxford University Press.