Minds & Machines (2013) 23:255–258
DOI 10.1007/s11023-012-9293-z

BOOK REVIEW

Jon Williamson: In Defence of Objective Bayesianism

Oxford University Press, Oxford, 2010, vi+185, $85.00 (hard), ISBN 978-0-19-922800-3

Hykel Hosni

Published online: 28 November 2012

The aim and spirit of Jon Williamson's *In Defence of Objective Bayesianism* are effectively summarised by the author as follows:

> This book is written in the belief that it is better to contribute to the struggle to state and defend the right position than to settle for a more easily defensible position that is only a part of the story. (p. 163)

According to Williamson, the "right position" is deeply rooted in common sense: "if someone's evidence leaves the truth or falsity of θ open, then she would be irrational to strongly believe θ or its negation". In spite of its intuitive appeal, a number of attempts at formalising this idea turned out to be unsatisfactory, with some leading to well-known paradoxes. The book under review takes up the two-fold challenge of articulating "the right position" in full formal detail and defending the ensuing principle of Maximum Entropy against a number of objections, old and new. The essence of Williamson's defence strategy is to entrench the principle in a foundationally robust framework—Objective Bayesian Epistemology (OBE).

The defence articulated throughout the book under review addresses objections of a practical and of a foundational kind. The former are mainly focused on the applicability of OBE methods in uncertain reasoning and decision-making. The latter question the status of OBE as a suitable normative framework for the representation of rational behaviour under uncertainty. The foundational criticisms addressed in this book may in turn be divided into anti-Bayesian and intra-Bayesian according to whether they take issue with Bayesianism *tout court* or with (some of) the specific principles which define OBE.

H. Hosni (✉)
Scuola Normale Superiore, Pisa, Italy
e-mail: h.hosni@gmail.com

H. Hosni
Centre for Philosophy of Natural and Social Science, London School of Economics, London, UK

My opinion is that *In Defence of Objective Bayesianism* largely succeeds in meeting the practical and the anti-Bayesian criticisms it takes into account. As to the intra-Bayesian objections, I think Williamson is perhaps *too defensive* in excluding interval-valued probabilities as suitable candidates for norms of rational belief. Just as Williamson shows that subjective Bayesianism relies on essentially the same arguments which support OBE, his justification for measuring uncertainty with single-valued probabilities doesn't appear to be so obviously inapplicable to the interval-valued case. Before I articulate this in some detail, let me briefly overview the contents of this particularly dense book.

Chapters 1–3 set up the stage by spelling out Williamson's characterisation of OBE. Given the author's overall aim it is not surprising that this part accounts for nearly one half of the book. In Chapter 1 Williamson lists the criticisms that he will address and delimits the scope of his investigation, namely a normative account of rational, evidence-based, propositional degrees of belief. The first part of Chapter 2 provides a compact but quite detailed historical account of Bayesianism "[f]rom Jakob Bernoulli to Edwin Jaynes". Whilst tracing the roots of OBE, Williamson points out how each of the three main interpretations of probability, namely the subjective, the physical and the logical, contributed—often by way of endorsing spectacularly radical positions—to the implicit consolidation of three norms of rational belief. The second part of Chapter 2 provides an explicit formulation of such norms: Probability, Calibration and Equivocation. Chapter 3, the longest in the book, is devoted to motivating them, leading to a full characterisation of OBE. Chapters 4–9 develop the necessary technical background to address specific counter-arguments to a number of foundational and practical criticisms. OBE turns out to be generally inconsistent with Bayesian updating, one of the cornerstones of standard Bayesian theory. Chapter 4 points out that when OBE and conditionalisation diverge, it is Bayesian conditionalisation that offers the weaker answer to the problem of formalising the dynamics of rational belief. This is Williamson's most radical departure from standard Bayesianism. Chapter 5 extends the characterising norms of OBE to predicate languages. This material provides a very useful addition to predicate probability logic, an area of mathematical logic for which there is currently an important revival of interest. Chapter 6 shows the applicability of OBE to Bayesian Nets. This chapter addresses one of the most prominent practical criticisms to OBE, namely that Maximum Entropy reasoning is too computationally demanding to be useful in artificial intelligence. Williamson counters this objection by showing how the framework of Objective Bayesian Nets (OBNET) leads to efficient Maxent reasoning and how this resulted in successful real-world implementations of OBNET. Two further applications of OBE are discussed in Chapters 7 and 8, namely as a semantics for probability logic and as a framework for evidence-based social decision making. Chapter 9 provides a formulation of OBE in the language of measure-theoretic probability. This richer language is argued to be too general for the purposes of OBE, which therefore is best framed in the structure provided by logical languages. Chapter 10 concludes by setting the agenda for future research in OBE.

It is apparent from this abridged list of themes that *In Defence of Objective Bayesianism* has a decidedly multidisciplinary appeal and should be of great interest to logicians, epistemologists, statisticians, computer-, cognitive- and social scientists. In addition to this, Williamson's norm-based characterisation of OBE puts

uncertain reasoning into a uniform foundational perspective. This is, in my opinion, the central contribution of the book.

De Finetti, easily the most radical exponent of the subjective approach to Bayesian theory, insisted that (finite) additivity is the only constraint that can be normatively imposed on a rational agent's degrees of belief. Additivity leads to coherence, which for him and other subjectivists, notably Savage, captures the logic of subjective rationality: Coherence prevents individuals from making blatant mistakes, but otherwise grants them maximum flexibility in the expression of those degrees of belief which—in their opinion—reflect most accurately their information, experience, confidence, personality, and so on. This form of subjectivism certainly does not require agents to disregard evidence—including statistical information—nor does it deny that vast interpersonal agreement can emerge among Bayesian reasoners. What de Finetti does deny though is that conformity to those OBE requirements should have *normative force*.

Let me recall the OBE norms in slightly more detail. The Probability Norm demands that rational agents' degrees of belief should conform to the laws of probability. The Calibration Norm requires that the subjective probabilities licensed by the Probability Norm should be based on the agent's qualitative and quantitative evidence, including causal and logical relations, known frequencies, or, if these make sense in the specific case of interest, single-case physical probabilities. Finally, the Equivocation Norm further refines the choice of degrees of belief by excluding extreme probability values unless these are being prescribed by the previous norms, and subject to this requirement, it constrains probabilities to be otherwise minimally prejudiced, or equivalently, maximally equivocal. Williamson shows (Chapter 3) that Probability, Calibration and Equivocation are individually justified by appealing to formal variations of essentially the same argument which involves the minimisation of a certain loss function. The gist of the argument can be described as follows. Provided that an individual is faced with a suitably defined choice problem, each of the three above norms can be justified by showing that contravening them would increase the agent's expectation of incurring a loss (possibly in the long run). Hence the framework subjective Bayesians have long used to articulate and defend *their* position, effectively provides a justification for *objective* Bayesianism. This leads to the ironic conclusion that if you think of yourself as a subjective Bayesian then you might have a very hard time objecting (consistently) to OBE.

I find Williamson's norm-based characterisation of OBE very compelling and I feel persuaded by most of his specific arguments in defence of it. A significant exception, as anticipated above, is his criticism of the interval-valued extension of standard Bayesianism, to which I now turn. Williamson considers various interpretations of interval-valued partial belief and concludes that single-valued OBE is overall better justified. After criticising some non-Bayesian approaches to probability intervals, notably the one based on Kyburg's *evidential probability*, Williamson addresses in Section 3.4.7 two main lines of intra-Bayesian criticism which are labelled "pragmatic and conceptual" (p. 70), respectively. As to the former, Williamson maintains that interval-valued degrees of belief are "harder to obtain and work with" than their point-valued counterparts. There is no doubt about this—complication is hardly a surprise when expressive power increases. Yet he goes on to suggest—and this brings us to his

"conceptual" criticism—that the cost of extending the formal apparatus of Bayesian theory to interval-valued probabilities may not be balanced by an adequate gain in expressive power. This is where I disagree with Williamson. His argument is in fact three-fold, but essentially pivots on the observation that "the interval approach [...] weakens the link between belief and decision" (p. 70). There are several levels at which this objection can be spelled out and, consequently, addressed, but I think it suffices to remain at the level of interpretation. Walley's approach to imprecise probabilities, which Williamson recalls at various points, is grounded in decision-theoretic principles of admissibility which can be construed as natural adaptations of de Finetti's coherence. This requires a crucial fine-tuning of the betting scenario. In particular, in the interval-valued case, gamblers cannot bet negative stakes whilst bookmakers can differentiate between buying and selling prices. Under this interpretation, the real-valued interval $[l, u]$ represents the supremum of the buying price and the infimum of the selling price, respectively, for an event of interest θ. The width of the interval can then be interpreted as expressing the 'second-order uncertainty' attitude of the agent with respect to their willingness to bet on θ. Two extreme choices are noteworthy, namely [0,1] representing 'unwillingness to bet' and $l = u$ capturing 'willingness to bet at fair odds'. Taking this 'attitude' into account in the quantification of a rational agent's uncertainty leads to an increase in expressive power that need not conflict with the norms of Bayesian epistemology. Indeed, building on this interpretation, Fedel et. al. (2011) have defended the normative force of a suitably adjusted criterion of coherence which accounts for Williamson's "second conceptual worry", namely that whilst the betting interpretation underlying interval-valued probability is naturally more realistic, it has dubious normative force.

It goes without saying that Williamson cannot be held responsible for not taking into account results which have appeared *after* the publication of his book. The point of this 'counter-objection from the future' is therefore limited to suggesting that his defence of OBE need not be incompatible, at least in principle, with (the Bayesian interpretation of) interval-valued norms of rational belief. Encouraging signals to the effect this might be the case also in practice come from the recent literature on norms of rational *decision* which develop the idea of weighing utility with interval-valued, rather than point-valued, measures of rational belief (see, e.g. Troffaes 2007). This clearly suggests that much could be gained by grounding OBE in an Imprecise Probability Norm. Whilst a substantial amount of work will be required to articulate a research programme along these lines, there is no doubt that Williamson's *In Defence of Objective Bayesianism* provides an indispensable reference for anyone with an interest in tackling it.

References

Fedel, M., Hosni, H., & Montagna, F. (2011). A logical characterization of coherence for imprecise probabilities. *International Journal of Approximate Reasoning, 52*(8), 1147–1170.

Troffaes, M. (2007). Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning, 45*(1), 17–29.