Michael G. Titelbaum
Quitting Certainties.
Oxford: Oxford University Press 2013.
368 pages
$75.00 (hbk: ISBN 978-0-19-965830-5)

Bayesian epistemology uses probability theory to formally model degrees of belief and provide coherence constraints on rational credence. This involves both a synchronic and a diachronic dimension. According to the synchronic constraint, for an agent's degrees of belief at a given time *t* to be rational and coherent, they should be probabilities $P_t()$ satisfying Kolmogorov's axioms:

Non-Negativity: $P_t(p) \geq 0$

Normality: If *p* is a logical truth then $P_t(p) = 1$

Finite Additivity: If *p* and *q* are mutually exclusive then $P_t(p \vee q) = P_t(p) + P_t(q)$

As for the diachronic constraint, it involves a probabilistic rule of inference given by a principle of conditionalization. The Bayesian formula for the conditional probability of some sentence *p* given another sentence *q* is given by the following ratio formula: $P_t(p \mid q) = P_t(p \,\&\, q) / P_t(q)$. For any time *t*, the certainty set C_t is defined as the set of all sentences *p* in the modelling language such that $P_t(p) = 1$. The set $C_k - C_j$ is thus the set of all certainties gained or what an agent learns between times *j* and *k*. Furthermore, we use the notation $[C_k - C_j]$ to represent a sentence that is the conjunction of the sentences in $C_k - C_j$ and if $C_k - C_j$ is empty then $[C_k - C_j]$ is the tautology T. Thus for times *j* and *k* such that $j < k$, an agent's rational degree of belief in sentence *p* at *k* is given by:

Conditionalization: $P_k(p) = P_j(p \mid [C_k - C_j])$

Despite its successes, there are issues with this standard Bayesian approach. The issue that is the focus of *Quitting Certainties* concerns the inability of the traditional Bayesian updating rule to properly model claims that move from certainty to uncertainty. This is because if upon conditionalization a sentence is assigned a probability of 1 it will subsequently always be assigned a probability of 1, even though it makes sense for an agent to be able to rationally reduce their probability assignment to a claim from 1 to less than 1 when moving from certainty at an earlier time to uncertainty at a later time. To deal with this issue, Titelbaum presents the Certainty-Loss Framework, a modified Bayesian approach that provides a way to model and accurately represent rational requirements on agents who undergo certainty loss. In doing so, he is able to provide a unified framework for dealing with scenarios in which agents lose certainty in a claim due to memory loss or because that claim is a context-sensitive one.

In Chapter 3 Titelbaum meticulously introduces the Certainty-Loss Framework (CLF), in particular its synchronic constraints. What is presented amounts to a system that gives a Kolmogorovian probability function for each time that is covered. Titelbaum however 'divides the labour' between *extrasystematic constraints* and *systematic constraints*. The job of extrasystematic constraints is basically to establish which claims an agent is certain of and which claims they are less-than-certain of at a given time along with any particular values the uncertain claims might take. This involves imposing a few certainty conditions. If an agent is certain of a claim *p* at time *t* or *p* is a claim deductively entailed by a certain claim then $P_t(p) = 1$. Otherwise *p* is assigned a specific value below 1 or a non-specific inequality of < 1. The systematic constraints, which are common to every model built using CLF, consist of a version of Finite Additivity and the Ratio Formula given above.

Given these certainty conditions, Subjective Finite Additivity guarantees that any unconditional credence function will be a Kolmogorovian probability function. Titelbaum provides some discussion regarding why he has divided his framework between extrasystematic and systematic constraints and this provides some insight into the well-thought-out and methodical nature of this approach. Chapter 4 looks at the application of CLF models to stories and discusses the normative implications of CLF and Chapter 5 looks at three common objections to Bayesian modelling frameworks: the problem of new theories, criticisms of the Ratio Formula, and logical omniscience.

Part III of the book covers the topic of memory loss, starting with Chapter 6, which introduces the central idea of Generalised Conditionalization (GC), CLF's diachronic updating rule that replaces traditional conditionalization. The Shangri La story provides a motivating example. In this story an agent is keeping track over three times of whether a coin toss came up heads (h). At t_0 before the toss their assignment is $P_0(h) = 0.5$. At t_1 after the toss $P_1(h) = 1$ before reaching t_2 at which point they can no longer remember with certainty that heads resulted due to some manipulation of their memory and therefore make the assignment $P_2(h) < 1$.

Given this move from certainty to less-than-certainty and the assignment $P_2(h) < 1$, adherence to the traditional Bayesian framework whereby sentences with probability 1 remain so gives a verdict of rationality violation because $P_2(h) = P_1(h \mid T) = P_1(h) = 1$. As Titelbaum argues though and one would be inclined to agree, agents do not violate the requirements of ideal rationality merely by forgetting the certainty of some sentence and it is perfectly reasonable to have $P_2(h) < 1$ in this story. Thus Generalised Conditionalization is brought in to deal with this issue:

> **Generalised Conditionalization (GC):** For any times t_j, t_k and any sentence p in the modelling language, if $P_j(\neg [C_k - C_j]) < 1$ and $P_k(\neg [C_j - C_k]) < 1$ then $P_j(p \mid [C_k - C_j]) = P_k(p \mid [C_j - C_k])$.

According to GC, conditionalization applies only when the set of certainties between times j and k increases. It also can be seen to have a certain symmetry; forward-temporal certainties gained can be seen as reverse-temporal certainties lost. Applying all of this to the Shangri La example, the problematic $P_2(h) = 1$ cannot be derived. Furthermore, what can be derived is: $P_0(h) = P_2(h) = 0.5$.

Chapter 6 finishes with a discussion of van Fraasen's Reflection Principle and chapter 7 contains an in-depth discussion justifying why GC should be adhered to. As evidenced by chapters 6 and 7, this alternative form of conditionalization proves to be an effective way to overcome the shortcomings of traditional conditionalization in dealing with cases of memory loss. It will also be successfully applied to cases of context sensitivity later on in the book. Given these successes, perhaps the CLF framework and GC in particular can be applied to other phenomena. One area that comes to mind is some type of probabilistic belief revision.

Part IV of the book covers the topic of context sensitivity. Because the truth-values of context-sensitive claims can change over time, an agent can go from certainty in a claim to less-than-certainty in that claim. For example, you might be certain at t_1 that 'it is now Monday' but lose certainty in that claim at t_2 because you have lost track of the time. As we have seen such change poses a problem for traditional conditionalization. Also, given that the status of some context-sensitive claim p can change between times t_1 and t_2, a conditionalization of the form $P_2(p \mid [C_1 - C_2]) = P_1(p \mid [C_2 - C_1])$ can be problematic. To deal with this issue, Titelbaum develops a method based on his Proper Expansion Principle that involves reducing a model whose language represents both context-sensitive and context-insensitive claims down to one whose language involves only context-insensitive claims, deriving some verdicts in that reduction and then carrying over those verdicts to the original model. Using this framework, given some context-sensitive claim q and context-insensitive claim p such that $P_2(p \equiv q) = 1$, the probability of q for t_2 can be calculated by first calculating $P_2(p \mid [C_1 - C_2]) = P_1(p \mid [C_2 - C_1])$ and then making a substitution to get $P_2(q \mid [C_1 - C_2]) = P_1(p \mid [C_2 - C_1])$.

Chapter 9 applies the full CLF to a range of stories, notably including the Sleeping Beauty Problem and Chapter 10 compares CLF to some other frameworks. Chapter 11 provides an interesting exploration of CLF, indifference principles and Everettian versions of quantum mechanics before the final two chapters wrap things up.

One point of particular interest from Chapter 12 is a comparative discussion of Jeffrey Conditionalization (JC) and an explanation of why the problematic stories looked at throughout the book would not have been successfully dealt with by using it. Nonetheless, JC does have potential benefits. Take the following two properties a framework might have:

(a) provides a way to successfully deal with certainty-loss and accurately model assignments of probability going from 1 to < 1

(b) can have diachronic changes in probabilities even if the certainty set does not change

Whilst CLF has the (a) property it does not have the (b) property. On the other hand, whilst JC has the (b) property it is the same as traditional conditionalization when it comes to certainties and thus does not have the (a) property. Might a hybrid that has both of these properties be devised? Perhaps such a framework could be used to model cases of gradual memory loss.

Overall this book is an impressive example of formal epistemology that tackles two significant challenges for traditional Bayesianism. The approach is methodical and apart from the core of these two primary tasks the book contains a number of valuable side insights. Finally, the outcome is a good example of how rather than seeking to adopt a fundamentally different alternative to a framework with certain issues (in this case Bayesianism), thinking carefully about the framework and making the right modifications to rectify the issues can be just the way to go.

Simon D'Alfonso
Independent Researcher