Michael G. Titelbaum, *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*, Oxford: Oxford University Press, 2013, xii + 345 pp., £40/$75 (Hardback), ISBN: 978-0-19-965830-5.

Quitting Certainties is an extremely ambitious treatise on Bayesian formal epistemology. The centrepiece is a 'modeling framework' which can be applied to non-ideal epistemic agents as well as to ideal agents. Unlike many traditional approaches to Bayesian epistemology, the framework allows for agents to go from certainty to less-than-certainty in a claim; Titelbaum accordingly calls it the 'Certainty-Loss Framework' (CLF). It is designed to model agents' degrees of belief in context-sensitive claims (such as the claim that it is now Monday), to allow for the forgetting of factual information (such as information about what was had for breakfast several years ago) and to allow for ignorance of logical truths (such as complicated mathematical theorems).

The content of CLF is inextricably bound up with the methodology behind its application. Throughout the book Titelbaum stresses the importance of taking an instrumentalist modelling-based approach to the Bayesian probability calculus. This approach has points of contact with the attitudes of de Finetti, Jeffrey, Garber and others, but Titelbaum explores it in unprecedented detail. He rejects interpretations according to which the Bayesian machinery exactly describes the properties or dispositions of some ideal agent, or according to which it describes normative conditions that – if met – guarantee that an agent is ideally rational. Instead, for Titelbaum the Bayesian probability calculus is a tool for building models of some particular epistemic 'story'; these models yield 'verdicts' which can then be used in assessing whether an agent's history of epistemic 'commitments' over the course of the story violates any (synchronic or diachronic) requirements of ideal rationality. These verdicts are model-specific; and crucially, a more detailed model may reveal requirements of ideal rationality which differ from the verdicts of a less detailed model.

Essentially, what Titelbaum's modelling procedure gives us is a set of necessary conditions for an agent in a scenario to count as ideally rational; he explicitly doesn't aim to give us a set of sufficient conditions. This is both a limitation of the approach and the source of some of its most interesting features. For example, Titelbaum's framework provides for an interesting treatment of the problem of logical omniscience. (This treatment, which develops ideas of Garber and others, involves a relaxation of the 'standard interpretation' of CLF; the standard interpretation does impose a logical omniscience requirement.) The basic idea behind the non-standard interpretation's approach is very simple: logical truths about which the agent is uncertain cannot feature in derivations of verdicts of the model. As a result, they cannot be the basis of judgments

that the agent has violated any requirement of ideal rationality. For example, on the non-standard interpretation an agent who fails to assign any degree of belief to the claim that P→(PvQ) does not violate any requirements of ideal rationality by having a higher degree of belief in P than in PvQ.

Although his modelling conception of the Bayesian calculus is rather distinctive, the basic updating machinery at the heart of Titelbaum's framework is familiar. He adopts a diachronic constraint which he calls 'GC', for 'generalized conditionalization'. Applying this constraint (under the assumption, about which Titelbaum is agnostic, that a unique degree of belief – or set thereof – is permissible in the light of any body of evidence) gives results that are essentially equivalent to the widespread 'ur-prior' approach which defines a full prior credence function and then conditionalizes this function on the certainties that an agent has at each time. Such an approach deals with memory loss in a straightforward way: if an agent has fewer certainties at a later time, then their degrees of belief at the later time are simply obtained by conditionalizing the agent's ur-prior on fewer propositions.

GC is not itself especially original, as Titelbaum admits; but the discussion which motivates it is one of the most subtle and interesting parts of the book. (Interestingly, the modelling methodology stressed elsewhere is very much in the background during this discussion.) Titelbaum grounds GC in a notion of 'suppositional consistency': roughly, being certain in A and supposing B should leave us with the same degrees of belief as supposing A and being certain in B. The idea is that the evidential relevance of a proposition to other propositions should be the same whether that proposition is supposed or believed with certainty.

The other key ingredient of CLF, central to Titelbaum's strategy for dealing with context-sensitivity, is a principle relating different models of the same situation. The 'proper expansion principle' (PEP) tells us when moving from a coarse-grained model of an epistemic story to a more fine-grained model is liable to result in different verdicts about ideal rationality. For example, if a coarse-grained model fails to represent some relevant information which the agent gains, then it will erroneously count the agent as violating requirements of ideal rationality if the agent changes their opinions in the light of that evidence. A finer-grained model representing the additional information will not generate this problematic verdict. Where PEP is satisfied – which is to say, when the agent at all times can provide a context-insensitive truth-value-equivalent in a coarse-grained model C for each of the context-sensitive claims in a finer-grained model F – then the verdicts of C can be safely carried over to F. Where PEP is not satisfied, then we need to use a model more detailed than C in order to obtain an accurate picture of ideal rationality's requirements.

A modelling framework like CLF is only as good as the models which can be built from it; so the final third of *Quitting Certainties* is taken up with analyses of controversial cases from the recent formal epistemology literature. Titelbaum develops, *inter alia*, models of Bas van Fraassen's Judy Benjamin problem, of Frank Arntzenius' Shangri La puzzle, of John Collins' prisoner example, and of Sarah Moss' mermaid problem; but the most detailed treatments are reserved for the notorious Sleeping Beauty problem and for the related problem of making sense of confirmation in the context of Everettian (many-worlds) quantum mechanics.

The puzzle of Sleeping Beauty emerged from the game theory literature in the mid-1990s and was popularized among philosophers by Adam Elga (2000). We are asked to imagine that Sleeping Beauty is to be put to sleep on Sunday night and woken just on Monday (if a fair coin lands Heads) or on both Monday and Tuesday (if a fair coin lands Tails). If the coin lands Tails and Beauty is woken twice, her memory is wiped between awakenings. So all three epistemically possible scenarios for Beauty upon awakening (Monday & Heads; Monday & Tails; Tuesday & Tails) are subjectively indistinguishable. Beauty knows the plan. The question is: what credence should she have, upon awakening on Monday morning, that the coin lands Heads?

A puzzle arises because there are apparently compelling arguments for incompatible answers. It seems that Beauty should have credence 1/2 in Heads, because she knew the coin was fair, and she apparently has no relevant new information after awakening. But it also seems that she should have credence 1/3 in Heads, since if the experiment is repeated many times the ratio of Heads-awakenings to Tails-awakenings tends to 1:2.

Titelbaum uses his framework to argue for the 'thirder' position in Sleeping Beauty. He gives two routes to the answer 1/3 in the context of CLF, both of which make use of the Principal Principle relating chance and credence. The first route closely follows an argument from Elga (2000); it involves both a Sunday evening/Monday morning model and a Monday morning/Monday evening model (representing a scenario in which part-way through Monday Beauty is told what day it is). Although the argument based on these two models (if sound) refutes the halfer position, to secure the thirder view the argument requires a controversial principle of indifference. (We might also worry about the way in which it combines the results of two distinct models.) So Titelbaum presents a further argument for the thirder solution which makes use of a single, more complex, model. This model applies not to Sleeping Beauty directly, but to a variant case called 'Technicolor Beauty' in which Beauty will see a single coloured paper – either red or blue – on each awakening.

I find both Elga's argument and Titelbaum's Technicolor Beauty argument for the thirder solution convincing. But of equal interest are Titelbaum's diagnosis of why

Sleeping Beauty cannot be solved directly via a single CLF model, and his claim that the Technicolor Beauty argument dispenses with the need for Elga's controversial restricted principle of indifference. I'll take these in turn.

Why can't Sleeping Beauty be solved with a single CLF model? According to Titelbaum, it's because on awakening there are no context-insensitive claims that Beauty is certain have the same truth-value as the context-sensitive claims which are evidentially relevant in the puzzle, e.g. 'Today is Monday'. By introducing the coloured papers, Titelbaum provides Beauty with context-insensitive truth-value-equivalents for these claims: e.g. 'the red paper day is Monday'. PEP then licenses us to extend verdicts from the coarser-grained model, which involves only context-insensitive claims, over to the finer-grained model, which includes the context-sensitive claims of interest. On this analysis, it is Beauty's lack of context-insensitive uniquely identifying descriptions for her awakenings which gives rise to the enduring controversies over Sleeping Beauty.

How can the Technicolor Beauty case dispense with the need for an indifference principle? Essentially, it makes use of a further chance event, with 50-50 odds, which is independent of the original coin toss that is the subject of Sleeping Beauty: a second application of the Principal Principle then replaces Elga's indifference principle. Titelbaum shows that some halfers are committed to shifts in their credences about the result of the second coin toss, even though that toss is physically completely independent of the first. This is a powerful trick, and it provides a vivid reminder of the strength of the constraints which the Principal Principle places on rational credence.

After defusing worries that the Technicolor strategy will generalize to situations where indifference is inappropriate, Titelbaum goes on to apply it to confirmation in Everettian quantum mechanics. The upshot of this discussion is that CLF is compatible with the Born rule for assigning probabilities in quantum mechanics, and that (assuming we can provide a suitable metaphysical picture according to which Everettian agents have degree-of-belief-like attitudes to particular outcomes) CLF leads to no troublesome 'automatic confirmation' of Everettian QM.

I found the discussion of the automatic confirmation problem to be less well-developed than other parts of the book. It is brief, and appeals to a model (UI) of an Everettian confirmation scenario in which there is no automatic confirmation effect. However, a structurally similar model can be constructed in the context of Technicolor Beauty which would then yield the halfer response. Titelbaum has the resources to say why that model delivers the wrong verdict about Technicolor Beauty, but he doesn't do so. He accordingly gives no explicit diagnosis of where the analogy between Technicolor Beauty and Everettian confirmation breaks down.

Titelbaum's book will be required reading for formal epistemologists. It makes a substantial contribution to the debate on updating degrees of belief in context-sensitive claims, it provides a distinctive and fruitful perspective on the methodology of Bayesian modelling, and it contains several other rich discussions which constraints of space have prevented me from describing. Moreover, it is a model of methodological clarity, and the elegant style makes it a genuine pleasure to read.

References

Elga, A. (2000). 'Self-Locating Belief and the Sleeping Beauty Problem', *Analysis* 60(2): 143-147.