

Titelbaum, M. *Quitting Certainties: A Bayesian Framework Modelling Degrees of Belief*, Oxford University Press, 2013, 345pp., £40 (hbk), ISBN – 978-0-19-965830-5

Reviewed by Martin Smith, University of Glasgow

Belief, according to Bayesians, comes in degrees. Furthermore, belief comes in degrees that can be represented by real numbers in the unit interval with 1 representing certainty. With the stage set in this way, Bayesians go on to offer a number of well known formal constraints prescribing how one's degrees of belief should be rationally managed. In *Quitting Certainties* Michael Titelbaum develops what he describes as a 'Bayesian' framework modelling degrees of belief. Titelbaum, though, is not orthodox Bayesian. His framework – which he dubs the *Certainty Loss Framework* – seeks to improve upon orthodox Bayesianism in a number of respects. I think that it does. As the name of the framework (and indeed the title of the book) suggests, its primary selling point is that it allows one to rationally lose confidence in claims of which one was previously certain.

Orthodox Bayesians lay down two formal constraints for the rational management of degrees of belief. One of these is a synchronic constraint that prescribes how one's degrees of belief should relate to one another at a given time and the other is a diachronic constraint that prescribes how one's degrees of belief should evolve over time. Let  $D$  be one's degree of belief function.  $D$  will be defined over a set of sentences, closed under the truth functional sentential operations  $\wedge$ ,  $\vee$  and  $\sim$ , and will take each of the sentences in this set to a real number.

According to orthodox Bayesians, if one is rational then one's degree of belief function must always conform to Kolmogorov's three probability axioms, *Normalisation*, *Non-negativity* and *Finite Additivity*:

For any sentences  $P$ ,  $Q$ ,

[N] If  $P$  is a logical truth,  $D(P) = 1$

[NN]  $D(P) \geq 0$

[FA] If  $P$  and  $Q$  are logically incompatible then  $D(P \vee Q) = D(P) + D(Q)$

According to orthodox Bayesians, if one is rational, then one's degree of belief function must always be a (Kolmogorovian) probability function. Two consequences of this are worth noting for what follows: First, if one is rational and  $D(P) = 1$  then  $D(P \wedge Q) = D(Q)$ . Second, if one is rational and  $P$  and  $Q$  are logically equivalent sentences, then  $D(P) = D(Q)$ .

Let  $D_t$  be one's degree of belief function at a time  $t$ ,  $D_u$  be one's degree of belief function at a later time  $u$  and  $L$  be the conjunction of sentences that one learns between  $t$  and  $u$ . According to orthodox Bayesians, if one is rational then the two degree of belief functions must conform to the principle of *Conditionalsation*:

For any sentence  $P$ ,

[CON] Provided that  $D_t(L) > 0$ ,  $D_u(P) = D_t(P | L)$

According to orthodox Bayesians, if one is rational, then one's degree of belief in  $P$  at  $u$  –  $D_u(P)$  – must be equal to one's degree of belief in  $P$  at  $t$ , conditional upon the conjunction  $L$  of everything that one learns between  $t$  and  $u$  –  $D_t(P|L)$ . Conditional degrees of belief are, in turn, taken to be defined by the ratio formula –  $D_t(P|L)$  is defined as  $D_t(P \wedge L)/D_t(L)$  if  $D_t(L) > 0$  and is undefined otherwise.

The orthodox Bayesian constraints force one to become certain of any proposition that one learns. Provided  $D_t(L) > 0$ ,  $D_u(L) = D_t(L \wedge L)/D_t(L)$ . Since  $L \wedge L$  and  $L$  are logically equivalent, if one is rational then  $D_t(L \wedge L) = D_t(L)$  in which case  $D_t(L \wedge L)/D_t(L) = D_t(L)/D_t(L) = 1$ . On the orthodox Bayesian picture, we have it that (A) *if one is rational, then whenever one changes one's degree of belief in a claim, there must be some claim of which one becomes certain*. According to (A) all rational changes in one's degrees of belief must be accompanied by the acquisition of certainties. Furthermore, once one does become certain of a claim, orthodox Bayesian constraints leave no room for one's degree of belief in that claim to ever be lowered again. If  $D_t(L) > 0$  then  $D_u(P) = D_t(P \wedge L)/D_t(L)$ . If one is rational and  $D_t(P) = 1$  then  $D_t(P \wedge L) = D_t(L)$  in which case  $D_t(P \wedge L)/D_t(L) = D_t(L)/D_t(L) = 1$ . We have it that (B) *if one is rational then, once one is certain of a claim, one must not change one's degree of belief in that claim*.

Taken together (A) and (B) seem to mandate a sort of dogmatism – a picture on which any changes in one's degrees of belief oblige one to acquire certainties and to cling on to those certainties come what may. Richard Jeffrey famously argued that learning need not always involve the acquisition of certainties. As such, he replaced Conditionalisation with a more relaxed constraint – which has come to be known as 'Jeffrey Conditionalisation' – that allows us to escape from (A) (see Jeffrey, 1965, chap. 11). Jeffrey's framework, though, retains a commitment to (B) – if one does become certain of a claim, Jeffrey conditionalisation leaves no room for one's degree of belief in that claim to ever be lowered again<sup>1</sup>. Titelbaum's framework, however, offers a way of escaping from (B).

One source of trouble for (B) is the possibility of *memory loss*. Suppose I decide one evening to roll a six sided die. Before I roll, my degree of belief that the die will come up 6 is  $1/6$ . I roll the die, it comes up 6 and I see that it does. At this point I become certain that the die came up 6. A year later, however, I've completely forgotten what I rolled that evening and my degree of belief that the die came up 6 is back to  $1/6$ . It's easy enough to imagine one's degrees of belief changing in this way, and such changes would seem to involve no irrationality. Forgetting things may be a failing of some sort, but it is not a *rational* failing. If (B) is true, though, then I must be guilty of some rational failing. This kind of change in my degrees of belief is not consistent with what the orthodox Bayesian framework prescribes.

Another source of trouble for (B) is the phenomenon of context sensitivity. If one is certain of a context sensitive claim – 'It's now April', 'It's currently raining' etc. – rationality clearly does not require that one remain certain of the claim for ever more – after all, it may change its truth value from one time to the next. Titelbaum proposes to replace Conditionalisation with two new constraints. One of these, which he terms 'Generalised Conditionalisation' (section 6.1.3), is designed specifically to handle cases involving memory loss while the other, which he terms the 'Proper Expansion Principle' (section 8.2), is designed to handle cases involving context sensitivity. The discussion of Generalised Conditionalisation and memory loss takes place in chapters 6 and 7. It is this discussion that I will focus on here.

Let  $C_x$  be one's *certainty set* at time  $x$  – the set of sentences of which one is certain or committed to being certain at time  $x$ . Let  $C_x - C_y$  be a set containing all of the sentences that are in  $C_x$  but not in  $C_y$ . Finally, let  $\langle \rangle$  be a function that, when applied to set of sentences, generates a conjunction of those sentences – so  $\langle \{P, Q, R\} \rangle = P \wedge Q \wedge R$ . Let  $\langle \rangle$  generate a logical truth when applied to the empty set. With these definitions in mind, the Conditionalisation constraint could be rephrased as follows:

$$[\text{CON}] \quad \text{Provided that } D_t(\langle C_u - C_t \rangle) > 0, D_u(P) = D_t(P | \langle C_u - C_t \rangle)$$

According to Conditionalisation, if one is rational, one's new degree of belief in  $P$  at a later time must be equal to one's degree of belief in  $P$  at an earlier time conditional upon all of the new certainties gained since then.

Titelbaum's Generalised Conditionalisation constraint is as follows:

$$[\text{GC}] \quad \text{Provided that } D_t(\langle C_u - C_t \rangle) > 0 \text{ and } D_u(\langle C_t - C_u \rangle) > 0, D_u(P | \langle C_t - C_u \rangle) = D_t(P | \langle C_u - C_t \rangle)$$

Generalised Conditionalisation has a pleasing symmetry to it – according to this constraint, if one is rational then one's degree of belief in  $P$  at  $u$ , conditional upon all of the certainties lost since  $t$ , must be equal to one's degree of belief in  $P$  at  $t$ , conditional upon all of the certainties gained before  $u$ .

If one only acquires certainties between times  $t$  and  $u$  then GC reduces to CON – that is, the two constraints will offer exactly the same prescriptions. If  $C_t - C_u$  is empty then GC becomes this: Provided that  $D_t(\langle C_u - C_t \rangle) > 0$  and  $D_t(T) > 0$ ,  $D_u(P | T) = D_t(P | \langle C_u - C_t \rangle)$  for some logical truth  $T$ . If one's degrees of belief conform to the probability axioms then  $D_t(T) = 1$  and  $D_u(P | T) = D_u(P)$  in which case this just becomes CON: Provided that  $D_t(\langle C_u - C_t \rangle) > 0$ ,  $D_u(P) = D_t(P | \langle C_u - C_t \rangle)$ . It is in this sense that GC represents a *generalisation* of CON. But if one *loses* certainties between times  $t$  and  $u$ , the prescriptions of GC and CON diverge.

Consider again the die case described above. Let time  $t_1$  be the time before I roll the die,  $t_2$  be the time immediately after I roll the die and  $t_3$  be the time a year later. Let  $P$  be the sentence that the die came up six on the particular night in question. We have it that  $D_{t_1}(P) = 1/6$ . Between  $t_1$  and  $t_2$  I learn that  $P$  is true – and this is the only change to my certainty set. Thus, GC, like CON, prescribes that  $D_{t_2}(P) = 1$ . What about  $D_{t_3}$ ? Between  $t_2$  and  $t_3$  I effectively lose the very certainty that I acquired between  $t_1$  and  $t_2$ . That is, my certainty set at  $t_3$  is equal to my certainty set at  $t_1$  –  $C_{t_3} - C_{t_1}$  and  $C_{t_1} - C_{t_3}$  are both empty. Thus, GC prescribes that  $D_{t_3}(P) = D_{t_1}(P) = 1/6$ , just like in the story. We are, of course, simplifying matters somewhat here. More realistically, I will have acquired all sorts of new certainties in the year between  $t_2$  and  $t_3$  – just none that are *relevant* to  $P$ . If we suppose that  $C_{t_3}$  is larger than  $C_{t_1}$  (and that  $D_{t_1}(\langle C_{t_3} - C_{t_1} \rangle) > 0$ ) then what GC prescribes is that  $D_{t_3}(P) = D_{t_1}(P | \langle C_{t_3} - C_{t_1} \rangle)$ . But, since  $\langle C_{t_3} - C_{t_1} \rangle$  is irrelevant to  $P$ , it's plausible that  $D_{t_1}(P | \langle C_{t_3} - C_{t_1} \rangle) = 1/6$ , in which case we still have the prescription that  $D_{t_3}(P) = 1/6$ . More realistically still, I will have acquired *and lost* all sorts of certainties in the year between  $t_2$  and  $t_3$ , but none that are relevant to  $P$ . If we suppose that  $C_{t_3}$  and  $C_{t_1}$  are partially overlapping sets (and that  $D_{t_1}(\langle C_{t_3} - C_{t_1} \rangle) > 0$  and  $D_{t_3}(\langle C_{t_1} - C_{t_3} \rangle) > 0$ ), then what GC prescribes is that  $D_{t_3}(P | \langle C_{t_1} - C_{t_3} \rangle) = D_{t_1}(P | \langle C_{t_3} - C_{t_1} \rangle)$ . But, since  $\langle C_{t_3} - C_{t_1} \rangle$  is irrelevant to  $P$ , it's plausible that  $D_{t_1}(P | \langle C_{t_3} - C_{t_1} \rangle) = 1/6$ , and, since  $\langle C_{t_1} - C_{t_3} \rangle$  is irrelevant to  $P$ , it's plausible that  $D_{t_3}(P | \langle C_{t_1} - C_{t_3} \rangle) = D_{t_3}(P)$  in which case we still have the prescription that  $D_{t_3}(P) = 1/6$ .

As long as we restrict attention to times  $t_1$ ,  $t_2$  and  $t_3$ , the prescriptions of Titelbaum's framework dovetail with the degree of belief changes that seem to be most natural in the story. It's not so clear, however, that the prescriptions of the framework are a good fit with what we might expect to be going on *in between* these times – in particular, in between  $t_2$  and  $t_3$ . One thing that we might observe about memory loss is that it would usually be a *gradual process*, and not a sudden change. Unless I suffer a bump on head or some such, there will be no instant at which I suddenly go from fully remembering rolling a 6 to having no memory of this event whatsoever – rather, the memory will slowly fade from  $t_2$  onwards. And this, presumably, will be reflected in how my degrees of belief change. That is, there will be no instant at which my degree of belief in the claim that the die came up 6 will change from 1 to  $1/6$  – rather it will start decreasing sometime after  $t_2$ , settling on  $1/6$  sometime before  $t_3$ . This pattern of change in my degrees of belief would seem to involve no irrationality. And yet, it is not clear that such change is consistent with what Titelbaum's framework prescribes.

As before, let  $P$  be the claim that the die came up 6 on the night in question. Let  $t_{2.1}$  be the time at which I first cease to be certain of  $P$  – the time at which  $P$  first drops out of my certainty set. If all of the other certainties that I've gained and lost since  $t_1$  are irrelevant to whether the die came up 6 then, by the above reasoning, GC prescribes that  $D_{t_{2.1}}(P) = 1/6$ . According to GC, my degree of belief in  $P$  should plunge, at  $t_{2.1}$ , from 1 to  $1/6$ . If I only become *slightly* less confident of  $P$  at  $t_{2.1}$  then Titelbaum's framework will predict that I am irrational. This prediction seems wrong.

One possible comeback to this is to argue that I will have other relevant claims in my certainty set at  $t_{2.1}$  that were not in my certainty set at  $t_1$  – perhaps the claim that (Q) I *seem* to remember rolling a 6, or some such. It is not clear that memory loss has to work in this way – that is, it's not clear that near-perfect confidence in a memory must be accompanied by perfect confidence in a *seeming* memory. But suppose we grant that  $Q$  is part of my certainty set at  $t_{2.1}$  and that  $D_{t_{2.1}}(P)$  is equal to, say, 0.99. Let  $t_{2.2}$  be the time at which  $Q$  first drops out of my certainty set. If my set of relevant certainties is now the same as it was at  $t_1$  then, by the above reasoning, GC prescribes that  $D_{t_{2.2}}(P) = 1/6$ . According to GC, my degree of belief in  $P$  should plunge, at  $t_{2.2}$ , from 0.99 to  $1/6$ . If I only become *slightly* less confident of  $P$  at  $t_{2.2}$  then Titelbaum's framework will predict that I am irrational. Once again, the prediction seems wrong.

We could argue of course that there are still relevant claims in my certainty set at  $t_{2.2}$  that were not in my certainty set at  $t_1$  – perhaps the claim that I *seem to seem* to remember rolling a 6. But then we might shift attention to time  $t_{2.3}$  at which this claim first ceases to be certain and so on. In order for GC to accommodate a *gradual* decrease in my degree of belief in  $P$ , we would need a large stock of claims that might slowly trickle from my certainty set between  $t_2$  and  $t_3$ . Titelbaum does briefly consider the possibility of a gradual reduction in confidence brought on by memory loss (section 12.2.1) – and suggests that it may well involve a gradual loss in underlying certainties. But the idea that every diminution in my confidence of a claim like  $P$ , no matter how slight, must be accompanied by a loss of certainties seems difficult to accept. I noted above that, according to Richard Jeffrey, learning need not always require the acquisition of certainties. The present point is, in a way, just the flipside of this: Forgetting need not always involve the *loss* of certainties.

The present problem needn't constitute an *objection* to Titelbaum's framework per se. Titelbaum is careful not to claim that the Certainty Loss Framework has a universal applicability –

indeed, he concedes that there will be certain situations in which its predictions do not represent genuine requirements of rationality (chapter 5). Perhaps cases of gradual memory loss will be amongst these situations. This would, I think, be a significant limitation – but it would not, in and of itself, threaten the application of the framework to situations like the original die case, in which we have just a few well selected ‘snapshots’ of a subject’s changing degrees of belief.

The problem that Titelbaum’s framework encounters with cases of gradual memory loss stems, in a way, from the fact that the framework continues to validate something close to thesis (A) above. On Titelbaum’s framework, it’s not true that all rational changes in degrees of belief must be accompanied by the acquisition of certainties, but it is true that all rational changes in degrees of belief must be accompanied by *changes in one’s certainty set* – by the acquisition *or* loss of certainties. We might call this (A\*). Jeffrey Conditionalisation avoids (A\*) and, as such, might provide a better way of modelling cases of gradual memory loss (Titelbaum floats this suggestion in section 12.2.1). Interestingly, though, this will only work if we begin as being somewhat less than certain of our memories. Jeffrey Conditionalisation, as noted above, retains a commitment to (B) – thus, if I start off certain that the die came up 6, Jeffrey Conditionalisation provides no way in which my degree of belief in this claim might ever be lowered. To do justice to cases of gradual memory loss that begin with certainty it may be necessary to employ a formal framework that dispenses with *both* (A\*) and (B). As far as I’m aware, no such framework has to date been developed.

I have focussed here on just one aspect of *Quitting Certainties* – namely, the GC constraint and the way in which it is able to deal with cases of memory loss. As I mentioned above, Titelbaum proposes another constraint on degrees of belief – the Proper Expansion Principle – which he uses (in conjunction with GC) to treat cases involving context sensitivity. The discussion of context sensitivity, in chapters 8-11, is intriguing and thought provoking – and includes an insightful discussion of the Sleeping Beauty problem (chapter 9) and an interesting exploration of quantum probabilities on an Everettian or ‘many worlds’ interpretation of quantum mechanics (section 11.3).

Another feature of Titelbaum’s book that is well worth highlighting is just how carefully and methodically he sets up the Certainty Loss Framework. He is, for instance, very careful to distinguish between the formal framework itself and the formal models that it generates, he is very careful to distinguish aspects of formal models and aspects of the informal real world situations that are being modelled and he is very careful in prescribing how rationality verdicts about situations might be read off formal results. When it comes to such things, he is a lot more careful and methodical than many Bayesians have been and, indeed, a lot more careful and methodical than I have been in this review. Irrespective of what one makes of the formal framework that Titelbaum defends, this book contains many methodological insights about the role and use of formal frameworks in general. Much of the discussion of general methodological issues takes place in chapters 2-5 and chapter 13 (see also section 7.5). Chapter 1 provides a brief introduction and preview of the book and chapter 12 mentions a few topics that don’t naturally fit into the discussion elsewhere – topics such as Dutch Book arguments, Jeffrey Conditionalisation and epistemic defeat. I would strongly recommend this book to anyone with an interest in formal epistemology – both for the innovative views that it contains and for the exemplary way in which Titelbaum goes about setting them out and defending them.

## ACKNOWLEDGEMENTS

Thanks to Philip Ebert for helpful comments on an earlier draft.

## REFERENCES

Jeffrey, R. *The Logic of Decision*, New York: McGraw-Hill, 1965

---

<sup>i</sup> Let  $D_t$  be one's degree of belief function at a time  $t$ ,  $D_u$  be one's degree of belief function at a later time  $u$  and suppose that the effect of whatever I learn between  $t$  and  $u$  is to alter the way my degrees of belief are distributed over a partition  $\{E_1, \dots, E_n\}$ . By 'Jeffrey conditionalisation' I have in mind the following constraint:

$$[JC] \quad \text{Provided that } D_t(E_i) > 0, \text{ for each } 1 \leq i \leq n, D_u(P) = \sum_i (D_t(P | E_i) \cdot D_u(E_i))$$

If  $D_t(E_i) > 0$ , for each  $1 \leq i \leq n$  and  $D_t(P) = 1$  then JC yields the result that  $D_u(P) = 1$ .