

Commentary on Brian T. McCann’s ‘Bayesian Updating’

Jochen Runde, Alberto Feduzi, and Laure Cabantous

2021-04-30 8:00:00

Editor’s Note

The following commentary was written in response to Brian T. McCann’s Using Bayesian Updating to Improve Decisions under Uncertainty published in California Management Review Volume 63 Issue 1 (Fall 2020). To read McCann’s direct response to Runde et al., please see Part 2.

In a recent article in this Journal, Brian McCann argues that making good decisions in the face of increasing uncertainty about the future requires thinking in probabilistic terms and acting as someone who “defines the set of possible outcomes [of some alternative] along with their associated values and probabilities . . . [and] then chooses the alternative that maximizes expected value” (p. 26). McCann observes that the quality of decisions made in this way depends on the accuracy of the probability estimates employed, and the purpose of his article is to advocate Bayesian updating as a means of increasing such accuracy.

McCann’s exposition is a model of clarity that owes much to his working example of a simple probability situation in which the decision maker is prepared to attach exact numerical values to their prior beliefs and beliefs about the strength of the evidence. Such values are a prerequisite for applying Bayes’ rule and we agree that Bayesian updating is appropriate, indeed follows as a matter of logic, when they are available. The question, however, is whether decision makers will always be prepared to assign such values, that is, numerically definite subjective probabilities, especially in situations of “deep uncertainty . . . ubiquitous in connected interdependent economies experiencing rapid technological change” (Teece et al. 2016 quoted by McCann 2020: 27).

We believe that the answer to this question is no and that this is why the familiar distinction between “risk” and “uncertainty” associated with the economists John Maynard Keynes and Frank Knight—between cases in which numerical probabilities can be determined and cases in which they cannot—has never quite gone away. Proponents of Bayesianism often respond to this distinction by pointing out that Keynes and Knight were writing before the advent of the “more expansive” subjective or personalist approach to probability favored by

McCann. But the fact that there exists a theoretical apparatus for reading subjective probabilities off preferences over lotteries as McCann describes does not mean that people actually do or, more importantly, always should, assign the precise numerical values required for Bayesian updating.

McCann does not comment on this issue and takes it for granted that a probabilistic thinker should be happy to trade in numerically definite probabilities: “Next, start assigning exact estimates to your degrees of belief” (p. 36), as he urges in the second of his simple steps towards becoming an explicit Bayesian. And we can see why many Bayesians might argue that any reluctance to do so would be unreasonable in light of the promise of possible “inaccuracies” in one’s priors being washed out by successive updating.

But all this begs the question of what probabilities becoming “more accurate” might mean, and whether one can talk sensibly of subjective probabilities becoming more accurate via Bayesian updating without there being “true” or objective probabilities for them to converge to. Some Bayesians, most famously the great Italian probabilist Bruno de Finetti, argue that this will never be the case, that there is no such thing as objective probability. We wouldn’t go this far and accept that there may be business-related situations in which there are underlying probabilities, frequencies perhaps, that are reasonably stable and to which subjective probabilities might converge as the evidence mounts (even if there are sometimes problems with such convergence being too slow when decisions have to be made quickly).

Further, we accept that some of the examples McCann provides of successful applications of Bayesian learning may be situations of this kind.

But equally it seems to us that there are many decision situations without objective probabilities to converge to, especially in emergent and fast-changing environments likely to give rise to deep uncertainty. While it would still be possible for someone to update their subjective probabilities in accordance with Bayes’ rule when presented with new evidence in such cases, there would be no guarantee that the updated probabilities would be any “more accurate” than their predecessors. In the absence of stable objective probabilities “out there” to be uncovered by Bayesian updating, subjective probabilities may be a poor guide even after repeated updating if the situation is one that keeps shifting.

These considerations bring us to a second and perhaps even more fundamental prerequisite of standard Bayesian updating, which is that the decision maker knows all of the possible outcomes relevant to some decision in advance. One of the consequences of this prerequisite is that it rules out cases in which the decision maker is surprised by an “unknown” outcome they hadn’t thought of before. Take the case of someone attempting to estimate the proportion of red balls by drawing from an urn they believe to contain only red and black balls, who, after drawing some red and black balls, proceeds to draw a yellow ball. Bayesian updating grinds to a halt at this point, because its machinery precludes adding new outcomes or updating a zero probability to a positive probability. A

Bayesian in this situation would have to start again, reformulate their outcome space, re-specify their priors, and resume sampling and updating. Furthermore, this process would have to be repeated every time a possible outcome not previously considered is encountered, and where none of the learning involved would be via updating priors using Bayes' rule.

The issue of unknown outcomes doesn't arise in McCann's example where the problem is the simple one of assigning a probability p to "entry will be profitable" (and accordingly $1 - p$ to "entry will not be profitable"). All possible outcomes are covered in this case, at least if we follow McCann in taking it that "entry will break even" is not considered a third possibility. But by proceeding in this way McCann provides a highly sanitized version of a problem that will often be considerably more complicated in practice, where decisionmakers will often want also to distinguish between different levels of profits and losses, and importantly, between the different states of the world in which these different levels of profits occur (e.g., a 10% return achieved by tying the business to a single customer vs. a 10% return achieved by serving a host of different customers). And here we quickly come up against the fundamental problem that decision makers are often unable to specify in advance all conceivable states of the world and therefore, often, all conceivable outcomes of their actions, something that leaves them open to surprises surely all the more likely in situations of deep uncertainty (in fact, if they know that their list of possible outcomes may be incomplete, they shouldn't be assigning classical probabilities that sum up to 1 in the first place).

In short, and as the COVID-19 pandemic and its myriad consequences again remind us, we live in a complex and emergent world. We accept that there are situations in which McCann's recommendations apply, namely where there are stable underlying frequencies that are relatively "closed" off from disturbing factors and the surprises these may generate. But outside of such situations, especially in the face of more radical "deep" uncertainty, it can be fundamentally misleading to adopt an approach that precludes surprise. Sticking within the framework of the Bayesian model does just this and can be doubly damaging for the false sense of confidence it may lend.

How then to proceed in light of these concerns, at least in situations that cannot be easily reduced into a form conducive to the application of Bayesian learning? There is much to be said here, but we will restrict ourselves to two suggestions.

The first is to pay more attention to the framing of decision problems, and then especially to the challenge of arriving at a reasonable idea of what the possible outcomes and often possible actions are in situations in which surprises are likely. Although often prior to assigning and updating probabilities, there are important questions of learning involved here too, if not of the kind associated with Bayesian updating. The philosophy of science and the psychology of reasoning provide rich sources of possible modes of reasoning, inference, and learning here, which are beginning to be adapted to arrive at methods to assist managers in the generation and exploration of possibility spaces and the uncovering of "unknown unknowns" before they can go on to becoming Black Swans (Feduzi & Runde,

2014; Feduzi et al., 2020).

The second is to go beyond the traditional cognitivist approach of focusing exclusively on what goes on in the mind of the decision maker and move towards a practice-based approach that pays more attention to the social and material aspects of decision making— the body as well as the mind, and the surrounding environment including tools (Cabantous & Gond, 2011; Cabantous & Gond, 2015; Cabantous et al., 2010). Indeed, McCann takes a step in this direction where he comments on the usefulness of spreadsheets, Bayesian calculators, and the like. But what we have in mind here is something more radical than this, namely that organizations pay more attention to creating tools, devising choice architectures, etc., in ways that facilitate sound judgments without making undue demands on individuals’ computational abilities. A good example of what we have in mind here is Vallée-Tourangeau et al.’s (2015) demonstration of how tools (in this case, a set of cards providing a malleable physical representation of the problem) can help agents update their probabilities in a manner consistent with, but *without* instruction and training in, Bayesian reasoning. In the same way, organizations may be able to alter their work environments by embedding in their routines and culture what are sometimes called “cognitive repairs” (that is, simple procedures such as the “Kokai watches” at Bridgestone and the “Five Whys” at Toyota) to facilitate uncovering unknowns without making significant additional cognitive demands on the individuals involved (Heath, Larrick and Klayman, 1998).

Read Next: McCann’s response to Runde et. al

References

1. Brian T. McCann, “Bayesian Updating to Improve Decisions under Uncertainty,” *California Management Review* 63/1 (2020): 26–40.
2. Laure Cabantous, and Jean-Pascal Gond, “The resistible rise of Bayesian thinking in management: Historical lessons from Decision Analysis,” *Journal of Management* 41/2 (2015): 441-470.
3. Laure Cabantous, and Jean-Pascal Gond, “Rational decision-making as a ‘performative praxis’: Explaining rationality’s éternel retour,” *Organization Science*, 22/3 (2011): 573-586.
4. Laure Cabantous, Jean-Pascal Gond, and Michael Johnson-Cramer, “Decision theory as practice: Crafting rationality in organization,” *Organization Studies*, 31/11 (2010): 1531-1566.
5. Alberto Feduzi and Jochen Runde, “Uncovering unknown unknowns: towards a Baconian approach to management decision-making,” *Organizational Behavior and Human Decision Processes*, 124/2 (2014): 268-283.
6. Alberto Feduzi, Phil Faulkner, Jochen Runde, Laure Cabantous, and Christoph Loch, “Heuristic methods for updating small world representations in strategic situations of Knightian uncertainty,” *Academy of Management Review* (2021) forthcoming. (DOI: 10.5465/amr.2018.0235)
7. David Teece, Margaret Peteraf, and Sohvi Leih, “Dynamic Capabilities and Organizational Agility: Risk, Uncertainty, and Strategy in the Innovation Economy,” *California Management Review*, 58/4 (Summer 2016):

13-35. 8. Gaëlle Vallée-Tourangeau, Marlène Abadie, and Frédéric Vallée-Tourangeau, “Interactivity fosters Bayesian reasoning without instruction,” *Journal of Experimental Psychology: General*, 144/3 (2015): 581–603. 9. Chip Heath, Richard P. Larrick, and Joshua Klayman, “Cognitive repairs: How organizational practices can compensate for individual shortcomings”, *research in Organizational Behavior*, Vol. 20 (1998): 1-37.