Is everyone Bayes? On the testable implications of Bayesian Fundamentalism

doi:10.1017/S0140525X11000409

Maarten Speekenbrink and David R. Shanks

Division of Psychology and Language Sciences, University College London, London WC1E 6BT, United Kingdom.

m.speekenbrink@ucl.ac.uk d.shanks@ucl.ac.uk

http://www.psychol.ucl.ac.uk/m.speekenbrink

http://www.psychol.ucl.ac.uk/david.shanks/Shanks.html

Abstract: A central claim of Jones & Love's (J&L's) article is that Bayesian Fundamentalism is empirically unconstrained. Unless constraints are placed on prior beliefs, likelihood, and utility functions, all behaviour – it is proposed – is consistent with Bayesian rationality. Although such claims are commonplace, their basis is rarely justified. We fill this gap by sketching a proof, and we discuss possible solutions that would make Bayesian approaches empirically interesting.

Although the authors are perhaps attacking a straw-man, we agree with many points raised in Jones & Love's (J&L's) critique of "Bayesian Fundamentalism." It is our objective here to strengthen their claim that Bayesian Fundamentalism is empirically unconstrained; although such claims are often made, their basis is not usually fleshed out in any detail. This is such a key part of the case that we sketch a proof and discuss possible solutions.

Without placing constraints on prior beliefs, likelihood, and utility functions, claims of Bayesian rationality are empirically empty: any behaviour is consistent with that of some rational Bayesian agent. To illustrate this point, consider a simple probability learning task in which a participant has two response options (e.g., press a left or a right button), only one of which will be rewarded. On each trial t, the participant gives a response $x_t = \{0,1\}$, and then observes the placement of the reward $y_t = \{0,1\}$, which is under control of the experimenter. The question is whether the assumption of Bayesian rationality places any restrictions on the response sequence for a given reward sequence.

In Bayesian inference, the prior distribution and likelihood (model of the task) assign a probability $P(y_t = S_j)$ to each possible reward sequence. Without further constraints, we can take this probability to be proportional to a value $v_j \geq 0$. After observing y_1, some of the rewarded sequences are impossible, and learning consists of setting the probability of these sequences to 0 and then renormalizing. For example, consider a task with three trials. The possible reward (and response) sequences are given in Table 1. Assume the sequence of rewards is $y = S_1$. After observing $y_1 = 0$, S_5 to S_8 are impossible and the posterior probabilities become $P(S_j|y_1) = v_j/\Sigma_k v_k$, for j, $k = 1, \ldots, 4$, and $P(S_j|y_1)=0$ for $j = 5, \ldots, 8$. After observing $y_2 = 0$, S_3 and S_4 are also impossible, and the posterior probabilities become $P(S_j|y_1) = v_j/\Sigma_j v_k$, for j, $k = 1, 2$, and $P(S_j|y_1) = 0$, for $j = 3, 4$. After observing $y_3 = 0$, only S_1 remains with a probability 1.

A rational Bayesian agent gives responses which maximise his or her subjective expected utility, conditional upon the previously observed rewards. For simplicity, assume the utility of a correct prediction is $u(y_t = x_t)=1$ and that of an incorrect prediction is $u(y_t \neq x_t) = 0$, so that the expected utilities correspond to the posterior predicted probabilities of the next reward. The crucial point is that in this general setup, we can always choose the values v_j to make any sequence of responses x_t conform to that of a maximizer of subjective expected utility. For example, suppose the sequence of rewards is S_1 and the sequence of responses is S_8. The first response x_1=1 implies that $v_1 + v_2 + v_3 + v_4 v_5 v_6 v_7 v_8$; the second response $x_2 1$ implies that $v_1 v_2 v_3 v_4$; the third response x_3 1 implies that $v_1 v_2$. One choice of values consistent with this is $v_j j$. For any response sequence, we can choose values which adhere to such implied inequalities, so behaviour is always consistent with a rational Bayesian agent. Although we have considered a rather simple situation with a small number of trials, this result generalizes readily to other sequential learning tasks such as category learning (for a related, more general and formal proof, see, e.g., Zambrano 2005). The problem becomes even more severe if we allow the utilities to depend on previous outcomes, which may not be entirely implausible (e.g., a third misprediction in a row may be more unpleasant than the first).

One may object that the particular method of Bayesian inference sketched here is implausible: Would someone really assign probabilities to all possible reward sequences? Maybe not explicitly, but in an abstract sense, this is what Bayesian modelling boils down to. Granted, the values assigned have been arbitrary, but that is exactly the point: Bayesian rationality is silent about the rationality of priors and likelihoods, yet some of these seem more rational than others. Thus, rationality hinges on more than adherence to Bayesian updating and utility maximization.

Is the claim of Bayesian inference and decision making always empirically empty? No. For instance, the assumption that rewards are *exchangeable* (that they can be reordered without affecting the probabilities) places equivalence restrictions on the values v such that, given a sufficient number of trials, some response sequences would violate utility maximization. Exchangeability is crucial to the convergence of posterior probabilities and the decisions based on them. Another option would be to let participants make multiple decisions while keeping their

Table 1 *(Speekenbrink & Shanks). Possible reward and response sequences (S_j) in a simple learning task with three trials (t)*

t	S_1	S_2	S_3	S_4	S_5	S_6	S_7	S_8
1	0	0	0	0	1	1	1	1
2	0	0	1	1	0	0	1	1
3	0	1	0	1	0	1	0	1

information base (posterior probabilities) constant, so that intransitive decisions become possible. More generally, testable conditions of Bayesian rationality can be found in the axioms of subjective expected utility theory (e.g., Savage 1954). Empirically meaningful claims of Bayesian rationality should minimally ensure the possibility that the data can falsify these axioms. Axiomatic tests are "model-free" in the sense that they do not rely on a particular choice of prior distribution and utility function. Such tests should be a first step in rational analysis; if the assumption of Bayesian rationality is not rejected, one can then look for priors and utilities which match the observed behaviour. Given rich-enough data, this search can be guided by conjoint measurement procedures (e.g., Wallsten 1971).

To conclude, while "Bayesian Fundamentalism" is generally unconstrained, by placing appropriate restrictions, the assumption of Bayesian rationality is subject to empirical testing and, when not rejected, can help guide model building.