

Rational Polarization: A Sceptical Puzzle for Truth-Seeking Bayesians

Anandi Hattiangadi & Joe Roussos

Abstract

Belief polarisation occurs when two agents' posterior beliefs move farther away from one another with respect to the same proposition or set of propositions. Polarisation has traditionally been regarded as a failure of rationality, e.g., the result of cognitive biases influencing the belief states of at least one of the polarising agents. However, polarisation is increasingly recognised as a potentially rational phenomenon that can result from: different evidence, different levels of trust, different priors, and more. Recently, Nielsen & Stewart (2020) have shown that in the extreme, two perfect Bayesian agents can receive the same increasing and infinite stream of non-misleading evidence and yet polarise, unless their priors are related in a way that satisfies strong and unmotivated constraints. What about convergence results, which purport to show that priors “wash out” and Bayesians converge on the truth in the long run? These results are subject to the same caveats which allow for rational polarisation: convergence is only “almost sure” and relies on strong and unmotivated assumptions about the agents' priors.

We argue that this raises a sceptical puzzle for truth-seeking Bayesians—those who aim for truth or accuracy in their beliefs—since it casts into doubt not only our confidence across the board, but the truth conduciveness of Conditionalization itself. To begin with, when two agents polarise with respect to a single proposition, one must move away from the truth, even if both are perfectly rational and confident in their assessment of the evidence. So, if you are one such agent, you cannot tell ‘from the inside’ whether you are in the good case—converging on the truth—or the bad case—converging on the false, for any proposition you believe or disbelieve. This casts the whole agenda into sceptical doubt.

Moreover, now that you know that rational polarisation is possible, you possess higher-order evidence that your reasoning process may be epistemically faulty, even in the absence of another agent with opposing beliefs. First, for every proposition that you believe or disbelieve there is a possible epistemic peer who has precisely the same evidence that you have, and who reasons precisely as you do, yet whose beliefs are polarised with respect to yours. This constitutes a peer disagreement case, writ large. Once again, this seems to cast the whole agenda into sceptical doubt. Second, whether you are in the good case or the bad case turns on facts about your priors that we may reasonably describe as a matter of luck. This seems to cast the truth-conduciveness of Conditionalization itself into doubt.

Can the truth-seeking Bayesian restore confidence in her beliefs and reasoning in the face of this sceptical doubt? Unfortunately, many anti-sceptical strategies fall short. We argue that Bayesian beliefs fail to meet credal analogues of sensitivity and safety conditions, and that Bayesians cannot defend either steadfastness or level-splitting in this context.

Nielsen, Michael & Rush T. Stewart. 2020. Persistent Disagreement and Polarization in a Bayesian Setting. *British Journal for the Philosophy of Science*, 72(1):51-78.