Bayesian Arguments for Weak Foundationalism

Gerhard Schurz, Düsseldorf

In this talk I give some simple arguments why there is an intrinsic need in Bayesian epistemology for a weak kind of foundationalism and for objective probabilities. My arguments are as follows:

1. Arguments based on (Jeffrey) conditionalization:

In Bayesian conditionalization of the simplest kind, we have two logically independent propositions A, B, and a prior probability distribution P_o over $\{\pm A \wedge \pm B\}$ ("±" for "negated or unnegated") at a time t_o. The probabilities are completely determined by the values of the nodes of the corresponding Bayes net $P_o(B|A)$, $P_o(B|\neg A)$, and $P_o(A)$. Already the orientation of the Bayes net indicates that we take "A" as epistemically prior to "B". But nothing in Bayesian epistemology makes this assumption explicit. At time t_1, a new experience is made, with the result that the initial probability of A at t_1, call it $P_{1\text{-}in}(A)$, is significantly greater than $P_o(A)$. Now, the probability $P_{1\text{-}in}(A)$ is incoherent with $P_o(B|A)$, $P_o(B|\neg A)$, and $P_o(B)$. So the question is: *how should one rationally change his or her probability values in order to make them coherent again.* We denote the new coherent probability distribution by P_1. Prima facie there are three possibilities: (i) change $P_o(B)$! ; (ii) change $P_o(B|A)$ and/or $P(B|\neg A)$, and (iii) change $P_{1\text{-}in}(A)$.

1.1 Bayesian conditionalization requires epistemic priority of evidence: What one never does in Bayesian conditionalization is to reset $P_{1\text{-}in}$ (A) to $P_1(A) := P_o(A)$. This would, for example, be the right kind of reaction if one is convinced of $P_o(B)$ to an extremely high degree, so that one concludes that the new experience leading to $P_{1\text{-}in}(A)$ was a wrongly interpreted or errorness experience. In Bayesian conditionalization, however, one takes the new initial probability value $P_{1\text{-}in}(A)$ for granted, i.e., one assumes $P_{1\text{-}in}(A) = P_1(A)$. This shows that Bayesian conditionalization presupposes a distinction between what counts as evidence – A – and what counts as predictive (singular) hypothesis – B. The probability values of evidence statements are immediately given as *inputs*; the probability values of hypotheses are calculated from the probabilities of evidence statements. This is a weak kinds of foundationalism: it implies an epistemic priority, though not an infallibility, of evidence statements.

1.2 Bayesian conditionalization requires objective and causally supported conditional probabilities: Even if we take it for granted that A counts as evidence, and hence set $P_1(A) := P_{1\text{-}in}(A)$, then there are still two possibilities left: we may change the unconditional probability of B, or we may change the conditional probability of B given A (and/or B given $\neg$A). In Bayesian conditionalization one assumes that $P_0(B|A) = P_o(B|A)$, and likewise for $\neg$A. In other words, the conditional probability of the prediction B given the evidence A is *stable* under new incoming evidence. But *why* should that hold? Note that p(B|A) is nothing but a quotient $P(A \wedge B) / P(A)$ of two unconditional probabilities. Moreover, I will give some examples where the change of the conditional probability will indeed be the more appropriate reaction. I claim that the natural reason why P(B|A) is considered as stable as against new incoming evidence is the fact that P(B|A) is considered as an objective, statistically interpreted probability value of the corresponding event-*types* p(Bx|Ax). But even this is not enough. It can be shown that in the usually Bayesian update rule, where one sets

$$P_1(B) = P_0(B|A) \,.\, P_1(A) + P_0(B|\neg A).P_1(\neg A) \quad \text{where}$$
$$P_1(B|\pm A) := P_o(B|\pm A)$$

it follows that the inverse conditional probability values *change* their value (at least in most cases) – that is, it will indeed hold that $P_1(A|B) \neq P_o(A|B)$ (and likewise for $\neg$B). So why is P(B|A) considered as invariant, but not P(A|B)? I claim that it is implicitly assumed that the direction from A to B reflects a directed (direct or indirect) *causal* influence from A to B. Only conditional probabilities reflecting causal influence are treated as stable in Bayesian conditionalization.

2. Bayesian arguments for non-circularity

Within the same framework of Bayesian conditionalization I will show that there cannot exist (completely) circular justification Circular justification means that, given some moderate probability value p(A), it is possible to justify B with help of A, leading to an increase of p(B), and then to justify A with help of the so justified B, leading to a further increase of p(A).

Consider the following graph

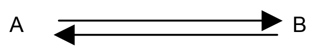

where the arrows mean high conditional probabilities. If p(A) increases, this will produce an increase of p(B) (given the conditional probability $p(B|\pm A)$ is held fixed.) However, this increase of p(B) *will not produce a further increase of p(A)*. This is a crucial difference to, for example, winner-take-all networks.

I will also show that Bayesian networks *do allow for partially circular justification*, which are displayed by the following graph:

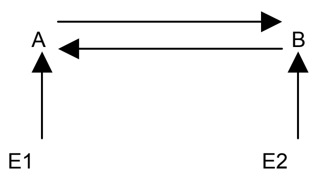

It can be shown that given an increase of P(E1) and of P(E2), the increase of P(E1) will have an effect on P(B), and the increase of P(E2) will have an effect on P(A), which is only possible because of the circular probabilistic support between A and B.

3. Bayesian arguments against extreme coherentism

Under extreme coherentism I understand the viewpoint that (i) the internal justification of belief set $B \subseteq L$ (L the formal language) is a function of our probability distribution over B (and over L), and that (ii) this probability distribution is chosen by the requirement of maximal internal coherence, where the internal coherence is a function of this probability distribution. Thus we choose our probability distribution for each $B \subseteq L$ in a way that the internal coherence of B becomes maximal. Then we choose that $B^* \subseteq L$ as our belief set which has highest coherence among all L-subsets .

If we map each atomic formula At into its negation and call this function f, we can construct for each probability function P an f-isomorphic probability distribution $P_{\neg}$ such that for each atomic formula At, $P_{\neg}(At) = P(\neg At)$. Moreover we get, for each belief set B, a corresponding belief set $B_{\neg}$. The unconditional and conditional P-probabilities over elements of B will have exactly the same values as the corresponding unconditional and conditional $P_{\neg}$-probabilities over the corresponding (subformula-negated) elements of $B_{\neg}$. Therefore the $P_{\neg}$-coherence of $B_{\neg}$ will be the same as the P-coherence of B. This means that by the method of extreme coherentism, we can construct for each maximally coherent belief set B many other belief sets B' differing from B in that for some atomic statements they assert the exact denial, which are equally internally coherent with respect to correspondingly modified probability functions.