In Defense of Reverend Bayes. Comment on Dennis

by
Charles E. McCulloch

Even as a dyed-in-the-wool frequentist I must cry foul with Brian Dennis' (BD's) rough treatment of Bayesians. Bayesians do test hypotheses and do calculate quantities very similar to (and in many ways *better* than) confidence intervals, p-values, and standard errors. To say that "Bayesian statistics provides conclusions in the face of incomplete data" is to obfuscate the fact that frequentists always use incomplete data. And BD's stream example in patently slanted towards a frequentist treatment.

I will elaborate on two of these points. Bayesians can calculate the probability that a hypothesis is false. This has a much more straightforward interpretation than the awkward definition of a frequentist p-value and is what many scientists would like to calculate.

BD's stream example is smack in the middle of the frequentist domain: it is easy to envision repeated sampling (even indefinitely) and there is only a single, easily estimated source of variability. But what is BD to do in devising a scientifically defensible conservation plan for an endangered species in the one location in which it remains? We might have some decent mortality data, only a little fecundity data (except perhaps in related species) and only expert opinion on survival if the habitat becomes fragmented. How do we combine these different sources of information and variability in a framework in which the frequentist paradigm is inappropriate? It is a hard problem but not one in which we should bury our heads in the sand.

Even in settings more amenable to frequentist inference the Bayesian paradigm can be attractive. In an area of interest to me, namely maximum likelihood estimation of variance components, how do we deal with setting confidence intervals? Even the large sample distribution of the estimates is an awkward mixture of point masses (as estimates hit the boundary of zero) and chi-square distributions with unknown mixing fractions. And this doesn't even start to deal with the small sample distribution, which cannot be written down in closed form. The Bayesian approach, together with modern computing power, straightforwardly handles this problem, even in small and moderate sized samples.

I must, however, agree with BD on the difficulties of specification of prior distributions. In my mind this is an Achilles' heel of the Bayesian approach. This is especially true in high-dimensional problems, in which the joint distribution of 25 parameters might be needed. As BD emphasizes, the specification is, in many cases, not innocuous and must be treated with care.

That said, prior distributions sometimes capture bonafide expert opinion, not just *personal* beliefs and are a synthesis of data and experience. BD also minimizes much of the current research on default, objective, and robust priors, all of which are intended to deal with the difficulty of personalistic priors. Also, ethical Bayesians invariably do sensitivity analyses to assess the influence of their prior specification.

I agree with BD to be extremely skeptical of the Bayesian approach in ecology. But let us not try to blame the ills of postmodern society and the ignorance of science by

the general populace on the Bayesians. And we should keep in mind situations where the Bayes paradigm is attractive and likely to be effective.

As regards BD's suggestion as to how to educate ecologists I say get with the program. Statistics minors at Cornell (a popular minor for ecologists) have, for over 20 years, been required to take a theory of probability and a theory of statistics course. To this is added a matrix algebra or modeling course as well as the usual yearlong data analysis and methods sequence. It makes an ecosystem of difference.