Science, Subjectivity and Software (Comments on the articles by Berger and Goldstein)

Anthony O'Hagan
University of Sheffield

May 4, 2006

Abstract

The dangerous heresy of so-called 'objective' Bayesian methods is again propounded by Berger. These comments are my attempt to save Bayesian statistics.

I have deliberately chosen rather dramatic, perhaps inflammatory, language in the above sentences. I do not expect all readers to view Berger's proposals in the same terms, but I hope that they will find my comments thought-provoking and constructive. It is undoubtedly true that the use of weakly informative prior distributions is both essential and valuable in practice. However, it is vitally important that their role is properly understood, instead of being grossly overstated.

My comments continue with some thoughts about Bayesian software that I hope are in tune with the agenda of the 'objective Bayesians', and certainly with the goal of spreading Bayesian methods more widely.

1 The role of weakly informative prior distributions

1.1 The 'O' word

I disagree with a great deal of what Berger says in his article, and I particularly dislike the use of the word 'objective'. It is simply a dangerously misleading term. I agree that it makes Bayesian methods seem more attractive if we call them 'objective', but what kind of justification is that? We could call them 'green', 'hospitable', 'small' or 'one-legged'; we could call *all* Bayesian methods 'objective' with as much justification; calling things 'objective' does not make them so, and is dishonest when they are not so.

It is true that 'objective' priors are constructed by following principles or algorithms, and that there may be no scope for subjective judgement in following them. But the choice of which principles or algorithms to follow is subjective.

It is a fact that different methods for construction lead to different priors. All are justified by their proponents according to principles that seem persuasive to them, but the multiplicity of such constructions means that there is no consensus on which are 'right' in any objective sense. To make claims of objectivity for any one approach is indefensible.

Scientists are taught that science is and must be objective; this unfortunately has served to hinder the progress of Bayesian analysis because it is perceived as inherently subjective. I understand the frustration which may have led fellow Bayesians to reach out to the physical scientists on this point, by trying to extract objective analyses and to promote them to scientists. However, I think this is misguided. It is more important to recognise the truth about science, which is that it is based on (a) careful separation of those things that are truly objective (i.e. the data) from subjective interpretation, plus (b) reasoned arguments and debate about interpretation, with the aim of achieving (c) agreement on interpretations and models. Bayesian analysis is perfectly in tune with this, but it does not belong in (a), which is where statisticians have resolutely tried to place their methods in the past. Bayesian analysis is about learning from data. This cannot be separated from interpretation of the data, which is subjective and brings in all the scientist's judgement and experience. Let us celebrate that process, recognise that it is necessarily subjective—albeit moderated by a rigorous framework for the debate—and not pretend that we can create objectivity merely by the use of words.

1.2 Philosophical viewpoints

Berger also has a section 1.2 with this title, in which he presents four philosophical positions. Unfortunately, I cannot agree with any of them, so let me present my own.

5. So-called 'objective' Bayesian analysis is a collection of convenient and useful techniques for approximating a genuine Bayesian analysis in appropriate practical situations.

The point is that the prior distributions that have been variously called 'objective', 'noninformative', 'reference', 'ignorance', 'weak' or 'default' priors *are useful*. I dislike several of these names, but 'objective' is the worst; I will call them *weakly informative* priors, because that is what I think describes them best. What I consider to be their value in Bayesian analysis is well known and alluded to by Goldstein in his article, but it is inexplicably ignored by Berger. So I will outline it here.

When prior information is weak, and the evidence from the data is relatively much stronger, then the data will dominate and the posterior distribution will depend only weakly on the prior distribution. In this situation, a weakly informative prior can be expected to give essentially the same posterior distribution as a more carefully considered prior distribution. The role of weakly informative priors is thus to provide approximations to a more meticulous Bayesian analysis.

This role is important, for two reasons. First, fully thought-through Bayesian analysis is a demanding task, so a quick and simple approximation is always welcome, particularly in view of the complexity and high-dimensionality of many modern Bayesian models. In the right circumstances, weakly informative priors provide very good approximations at much less effort than a full Bayesian analysis. (Goldstein's Bayes linear methods are another example of a technique that, in the right circumstances, can produce good approximations at vastly less effort.) The second reason why this is important is that the situation of weak prior information is one where it is particularly difficult to formulate a genuine prior distribution carefully. So it is especially convenient that this is when we may be able to get good answers with weakly informative priors.

The fact that there are several different weakly informative priors for a given problem, derived from different principles or algorithms, can be an advantage. From the philosophical viewpoint 5, we do not need to consider only one prior. We should always check the sensitivity of the posterior analysis to the prior distribution, in order to verify that the data are genuinely strong enough to override the choice of prior; having several possible weakly informative priors available is therefore a bonus.

For this reason, I use weakly informative priors liberally in my own Bayesian analyses. Often, there are just a few parameters about which prior information is substantial and worth formulating carefully; for the remainder, weakly informative priors suffice. But let me emphasise that I would *never* give to such analyses any of the interpretations of objectivity that Berger would apparently wish them to have. They are approximations to the analyses that I might be able to perform given more time and resources. I judge the approximations to be good enough to make further expenditure of time or resources not cost-effective. Everything we do in practice is an approximation in exactly this sense: there is nothing special about using weakly informative priors in this way.

2 Bayesian software and default analysis

Berger is right in saying that most statistics is not done by statisticians. They can do so because of the widespread availability of statistical methods programmed into standard software. Software makes it easy to do a wide range of statistical analyses, some quite sophisticated. Indeed, this is how almost all statistics is done in practice, whether by statisticians or others. The absence of user-friendly Bayesian software is a familiar complaint, and one that is certainly hindering the uptake of Bayesian methods at least as much as any concerns over subjectivity. In the application areas where I mainly work, there is almost no reluctance to embrace the use of expert knowledge, but great frustration over the lack of Bayesian software.

Of course, our community owes an enormous debt to the developers of BUGS and WinBUGS, the availability of which has had a profound influence on the progress of our subject and its penetration into so many application areas. But nobody can call WinBUGS user-friendly (I have even called it 'user-hostile' in my less positive moments!). Developers of major software products know that programming the underlying computations is a small part, and at least 80% of their resources must go into the user-interface that makes it as simple as possible for people to access those computations. The BUGS developers were never resourced at that level, and sensibly directed the bulk of their efforts into the creation of a package with very substantial computational power and functionality.

There is an old debate about the feasibility of Bayesian software. Frequentists can produce an analysis that purports to be applicable wherever a given model holds, regardless of prior information, whereas if we wish to create Bayesian software for that model we need to allow for all kinds of prior information. As a result, some have held the view that it will never be possible to create a Bayesian software package to rival the standard frequentist ones. To some extent this is true: standard packages trade on the lie that prior information does not matter. They offer the tempting delusion that statistics can be easy, and can be done with essentially no training. Unfortunately, 'objective' Bayesians are vying with them.

But just as I feel that weakly informative priors are useful—just not in the way that the 'objectivists' claim—I believe there is a role for Bayesian software in which weakly informative priors provide a kind of 'default' Bayesian analysis. The software that I envisage has the following characteristics.

- It tackles all the standard analyses that existing industry-standard packages do.
- It offers the user the option to employ weakly informative priors as a kind of 'default' option.
- It makes it clear to the user that 'default' here really means 'in default of other information that would make a useful contribution in addition to the data,' that better analyses are available to make use of genuine prior information, and that the 'default' analysis is at best an approximation to more considered, genuinely Bayesian, analysis.
- It makes the formulation of genuine prior distributions and subsequent analysis as straightforward as possible (e.g. by first offering suitable conjugate or conditionally conjugate priors).

There are numerous other desiderata, but these are the most important from the perspective of these comments. Such a Bayesian software package would make it easy for novices to begin to use Bayesian methods, and would provide a clear upgrade path from weakly informative priors where appropriate. I hope and believe that it would gradually draw many people into giving serious consideration to their prior information. However, it will also offer quick, good approximations when appropriate (making it clear when that is).

Given sufficient resources, this package could be developed *now*, building on WinBUGS. BUGS already does all the standard analyses with weakly informative priors, in a way that is so robust as to demand no user expertise.

The whole of the BUGS machinery for such analyses could be hidden from the user, providing instead an interface essentially like other packages. However, the user would have to explicitly choose 'default' priors each time, with alternative choices giving access instead to analyses based on genuine prior information.

I share the desire of 'objectivists' to attract more people to Bayesian analysis, but I do not wish to do it under false pretences.

3 A final plea

The search for 'objective' priors has already attracted more research and diverted more of our bright young Bayesian researchers than is warranted by their role as weakly informative priors, yielding approximations to genuine Bayesian analysis. I believe that to make claims of objectivity for these priors is disreputable; instead I would like our brightest researchers to be tackling problems that are honestly important for Bayesian analysis. Chief among these is how to formulate genuine prior information—particularly how to elicit it with the greatest possible fidelity and reliability from subject-matter experts. This has been a passion of mine for some time. The better we can understand how to formulate genuine prior information in complex problems, the easier it will be to become better Bayesians. Interested readers (and I sincerely hope there are many) can find some pointers to research and accepted practice in elicitation from the following recent book, and from the extensive bibliography therein.

O'Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J. Oakley, J. E. and Rakow, T. (2006). *Uncertain Judgments: Eliciting Experts' Probabilities*. Chichester: Wiley.