Bayesian Analysis (2006) **1**, Number 3, pp. 457–464

Rejoinder

James Berger*

My thanks to all the discussants for engaging in this interesting debate, and for the CMU Case Studies Workshop and *Bayesian Analysis* for making it happen. I should warn that I will engage in the spirit of the debate and take a more extreme objective Bayesian position than I perhaps believe. I'll first make a few comments about the companion paper by Michael Goldstein.

1 Comments on Michael Goldstein's paper

Michael makes a very nice, pleasantly positive, and practical defense of the subjectivist position. There is little he says that I disagree with, but I do have a somewhat different perspective on some issues based on my own experiences.

In *Applied Subjectivism*, Michael first visits a number of complex practical problems where subjective Bayes is argued to be needed. One example is analysis of complex computer models of processes which, coincidentally, is a problem with which I have also been seriously involved (Bayarri et al. (2002), Bayarri et al. (2005a), Bayarri et al. (2005b)). In our work we found that some subjective elicitation was needed, but we had to spend much more time dealing with objective Bayesian issues than with subjective issues. (The ratio of unknowns that we had to deal with objectively versus subjectively was consistently about 15 to 1).

In *Scientific Subjectivism*, Michael makes the case that the progress of science should be viewed in a subjective Bayesian sense; no argument here. But individual pieces of the process are often best stated in objective terms. Take the ESP example. It is very useful to separate out the prior probabilities of the hypotheses (i.e., individual opinions as to whether ESP exists or not) from what the data has to say (the Bayes factor of "Harry has ESP" to "Harry does not have ESP.") Being able to separate opinions from "what the data says" is important for the progress of science, and important in order for people to develop an understanding of data. Of course, this is a situation where a complete separation is not possible (since the Bayes factor can depend on parameter prior distributions), although objective Bayesians do strive to develop at least conventional Bayes factors which allow communication of what the data has to say. Going into this further would take the discussion too far afield, but a basic tenet of objective Bayesianism is that it is good to separate prior opinions from information in the data, to the extent possible.

The discussion of the court case is another example of this. My impression of the legal situation in the U.S. (the U.K. might well be different, of course) is that an expert witness cannot legally quote posterior probabilities of, say, guilt, but can only present

*Duke University and SAMSI, Durham, NC, http://www.stat.duke.edu/~berger

 ba0003

evidence such as Bayes factors. Thus it is imperative that objective Bayesian methods be utilized in court to communicate the information in data. (I am not saying that a scientifically established subjective distribution could not be used – e.g., a known distribution of a genetic trait in a population – but careful separation of prior opinion and information from data is needed.)

Next, Michael observes that full subjective Bayesian analysis via the usual approach is forbidding in its complexity, and then discusses the principles behind the Bayes linear approach, arguing that this approach is considerably more practical than standard full subjective Bayesian analysis. I have not tried Bayes linear analysis, and so do not know enough to discuss it intelligently, but it is interesting to note that Michael is not making a ringing endorsement of what is perceived as standard subjective Bayesian analysis.

Finally, in *Pragmatic subjectivism*, Michael points out the many uses of non-subjective prior distributions for a subjectivist. There is nothing I disagree with in his comments – indeed, my Section 2.4 makes many of the same points. In some sense, all I do is take this one step farther, arguing that one must know what one is doing in using non-subjective prior distributions, just as one must know what one is doing in using subjective prior distributions.

2 Response to J. Andrés Christen

I certainly agree that we should stop saying that subjective analysis and Bayesian analysis are synonyms. Of course, making major name changes is difficult, so I'd settle for Bayesian analysis being used to refer to the entire paradigm, with the adjectives *subjective* and *objective* being applied as desired.

3 Response to David Draper

I find myself at least in sympathy with all the general comments that David makes; I think we recognize the same issues, even if we have slightly different takes on how to deal with the issues. As to the specific comments:

- The pseudo-Bayes method of choosing "a uniform prior over a range that includes most of the mass of the likelihood function, but that does not extend too far," is not a method that I would expect to have good calibration properties. The devil is in the details: if one chooses a uniform prior too tightly concentrated around the likelihood function, one is close to just squaring the likelihood, which would have terrible calibration properties; on the other hand, if one chooses the uniform prior to be too dispersed, one is effectively using a uniform prior, which is known to have bad calibration properties in many situations. In a particular context, one might be able to study the situation and find the right balance between the two extremes, but that is my point: good objective Bayes analysis (or good calibration) is not guaranteed by just anything – one has to work to achieve it. Of course, once the work is done in a particular context, the objective Bayes problem is "solved" and the solution can henceforth be used for that context. (For the specific medical diagnosis problem in my paper, it is probably the case that uniform priors over the entire parameter space do give good frequentist calibration; in this example one would thus probably not need to worry about that particular extreme, but there are plenty of problems where both extremes are worrisome.)

- Objective Bayesian analysis in model choice is indeed much harder, because one can usually only utilize diffuse priors on parameters that are, in some sense, common between considered models. At the same time, one can argue that objective methodology is even more necessary here: if one is considering variable selection with 2^{60} models, it is obviously not possible to directly elicit 2^{60} different subjective priors. Also, in the model building phase of analysis, it is rarely worthwhile to spend a huge effort on prior elicitation for parameters in a model, because tomorrow the model might be discarded for something completely different. I shall resist the temptation of saying more, because model selection is a can of worms for both objectivists and subjectivists.

4 Response to Stephen Fienberg

Regarding the *History of Bayesian Thought*, I had not at all meant to suggest that Bayes or Laplace approached probability itself from an objective perspective. Indeed, I also view probability through the lens of degree of belief. It is objective Bayesian inference itself that I claimed had a long and illustrious history.

In Section 3.3 of my paper, I argued that one has to be careful in saying what the conclusion is from the *normative* perspective. There are a variety of different normative theories, and they do not all lead to the same conclusion. I do agree, of course, that it is more difficult to find axiom systems that lead to versions of objective Bayes than to find axiom systems that lead to versions of subjective Bayes.

Discussion of *objective scientists* often founders on the definition. When modeling data, scientists are more than happy to build in subjective knowledge, if it is the accepted knowledge in the field; to them, however, this is not being *subjective* – it is just being a good scientist. The story is far different when you get to aspects of the model about which there is no scientific knowledge or, at least, no scientific agreement. In my experience, scientists almost unanimously want to handle such aspects objectively.

The objective Bayesian search to me is not so much a search for the *holy grail* – a single grand overarching principle that will say what to do, but rather a search for, say, the *holy cross*, one sliver at a time (i.e., finding good objective Bayesian methodology in one specific situation after another). Of course, after enough "slivers" are done, general guidelines emerge, but these are not viewed as particularly holy.

The two examples Steve mentions are interesting in this regard. The Jeffreys prior for a covariance matrix has been known to be bad for roughly 30 years and, in the last 10 years, there has been considerable interest in finding good objective priors (because computation with good objective priors had become feasible). See Yang and Berger

(1994) and Berger et al. (2005) (in the references in my original paper) for discussion and other references. The example is also interesting because finding believable subjective priors for a covariance matrix is an incredibly daunting task (except when the covariance matrix is assumed to have a simplified form), so the example is actually almost a poster-child of the need for serious objective Bayesian analysis.

The second example mentioned is that of finding a good subjective prior for a large sparse contingency table. This is a problem that has not received much attention in the objective Bayesian literature, and is probably a great problem for study.

The conclusion of Steve's discussion is, I think, somewhat off-point. It is not what works best for me, or what works best for Steve, or what works best for *you* (given that you are still reading this article); I worry most about what works best for the 99% of statistical users who are not (and never will be) sophisticated Bayesians.

5 Response to Jay Kadane

I had not expected to hear a good new name for objective Bayes but I do quite like "interpersonal Bayes." It is probably not a suggestion that will catch on (unless Jay abandons the subjective viewpoint and devotes himself to the interpersonal approach), but the name does convey the notion that it is a Bayesian analysis that conveys useful information for everyone, and not just those who happen to trust the subjective Bayesian analyst.

Jay points out that "To admit that my model is personal means that I must persuade you of the reasonableness of my assumptions in order to convince you to continue reading the work." This sounds nice, and is certainly right in the abstract but, in practice, there are a number of difficulties:

- A good elicitation process often involves an enormous amount of detailed discussion and work between the subject matter specialist and statistician; effectively conveying all this in a report or paper is virtually never feasible. Indeed, most subjective Bayesian studies give only cursory explanations as to the process leading to the prior distributions.

- The prior distributions themselves are typically available for inspection but, absent a detailed understanding of how they arose, one usually can not do much more than simply trust that a sound elicitation was performed.

- In a truly complex study, properly digesting all subjective assumptions is almost overwhelmingly difficult, leaving one again in the situation of simply having to trust the analysts.

In contrast, when I look at a good objective Bayesian analysis, I can much more quickly judge whether the analysis is to be trusted or not, as I do not need to check the objective parts of the analysis that are based on what is known to be sound objective practice.

6 Response to Rob Kass

Rob has produced a wonderful set of questions, that everyone should try to answer for themselves (or study enough to figure out their answers). Here are mine.

1. Yes, if we are talking about procedures for routine use, and if good frequentist operating characteristics are compatible with good conditional performance (not always the case).

2. Yes: for a one-off analysis, Bayesian analysis can produce wonders.

3. I am not sure that anyone would answer negatively. Even the most extreme objective Bayesians do not discount the value of subjective Bayesian analysis in certain situations.

4. This relates to Philosophical Perspective 2 in my paper, which I basically agree with. The uncomfortable part of this to a Bayesian is the need to make the interpretation more context-dependent than we would like.

5. I would answer yes and I agree with the first part of the discussion, but view it as too limiting to say that the situation should be such that any reasonable prior belief would be overwhelmed by the data. This also relates to the previous question; I think that there is something special about serious objective Bayesian answers – they convey what you can learn from the data under an assumption of minimal prior information (not from *any* reasonable prior information). For instance, in a location parameter problem with one observation, the objective Bayesian answer is something with which I would be very comfortable, even though the answer is certainly not robust with respect to any reasonable choice of prior.

6. Yes, but I have nothing more to add after the last two answers.

7. No. I do not think the foundation of Bayesian inference is subjective, in the sense that subjectivism is why Bayesian inference is successful. I reiterate that I see very few pure subjective Bayesian analyses in practice.

8. Hmmmmm. I think the distinction between scientific inference and decision making is important, but I am not sure that it is central to the objective/subjective Bayesian debate. Problems of decision making will tend to involve more subjective analysis, while scientific inference will involve more objective analysis, but there will typically be a mix of each needed in either domain.

9. Hmmmmmmmmm. More m's because this is a really tough question that I had not thought about before. It is clear, I guess, that if the scientist is being consulted in a decision-making context, elicitation might well be useful. If the context, however, is simply advancement of science, the answer is far less obvious (see my rejoinder to Fienberg about objective scientists).

7 Response to Frank Lad

It is difficult to argue against Frank's introductory comments, in that it would probably be nice if the world could become the utopia he imagines. Alas, I do not see that as happening.

As to the technical questions:

- The first sentence of Frank's comment about the medical diagnosis problem is incorrect in that it refers to a "misleading promise." The word "misleading" should not be there, because the promise was realized – the frequentist question posed by Dr. Mossman was answered. Frank disagrees with the question, but that is an issue that I will leave him to argue with the psychiatrists. (It is not that the rest of his technical comment is irrelevant – far from it, there are serious issues being raised – it is just that we must be realistic about when and where these more delicate issues should enter our dialogue with others.)

- Frank's second point is very well taken. When one is in a scenario in which continual evaluation of modeling choices is possible, success in prediction is clearly the gold standard. Interestingly, however, the serious examples of this that I have seen (e.g., weather prediction) operate in a fashion that is not recognizable as any kind of probabilistic updating – the methodology that is used is often based on adhoc fiddling, done until the fiddling does better in prediction. This is thus another example where the philosophy of subjective Bayes is right, but the goal is usually achieved by something far different.

- Frank's last technical comment is not exactly what I would call a technical comment. Suffice it to say that I do not disagree with everything in the comment.

8 Response to Tony O'Hagan

As always, it's fun to interact with Tony, if only because he is not bashful in stating his opinions! Still, now I'm going to have to spend even more time trying to convince the world that Bayesianism is not a religion of fanatics; they will all quote the beginning of Tony's discussion in which he calls a very old and established variant of Bayesianism a "dangerous heresy".

By my "philosophical viewpoint 4" I more or less meant Tony's viewpoint 5; I took it for granted that the "learning from data" was meant in a Bayesian sense.

Tony's discussion about his common use of "weakly informative" priors in analysis, because their use will approximate his subjective Bayesian analysis (if only it could be done) is fine, but gets uncomfortably close to my comments on pseudo-Bayes procedures. Too often I see people pretending to be subjectivists, and then using "weakly informative" priors that the objective Bayesian community knows are terrible and will give ridiculous answers; subjectivism is then being used as a shield to hide ignorance of which "weakly informative" priors are appropriate, and which are not. There is often lip-service given to sensitivity analysis by varying the vague prior in such cases, but the variation in the prior that is done is often of an irrelevant sort. In my own more provocative moments, I claim that the only true subjectivists are the objective Bayesians, because they refuse to use subjectivism as a shield against criticism of sloppy pseudo-Bayesian practice.

I completely agree with Tony's comments about software. However, I would guess that 99% of the users will choose the "weakly informative" prior option. Since we know that there are terrible and good "weakly informative" priors, why shouldn't we then try to base such software on the good ones?

As to Tony's final plea, it's lucky that most bright young researchers don't listen to we old-timers who say "you should be doing what I think is important." That said, I'm an old-timer who also believes that not enough attention is being given to the methodology of elicitation.

9 Response to Larry Wasserman

The play was great, but I noticed it did not (alas) make the Academy Awards. Basically, I applaud the sanity check of asking how can we not care at all about the long-run performance of our procedures?

I do, however, have trouble with saying that good frequentist coverage is a necessary requirement. For instance, in the mentioned low count physics experiments, I am not convinced that one can simultaneously guarantee frequentist coverage while attaining good conditional performance. More directly, there is the large Gleser-Hwang class of problems (Gleser and Hwang (1987)) where *any* procedure that guarantees 95% frequentist coverage must, for some data, conclude "whole real line with confidence 0.95," a silly statement. So I would modify Larry's comments to say that frequentist coverage is a very desirable requirement, when attainable in a way that is conditionally sound. It is this sometimes unavoidable conflict between frequentist coverage and conditional sensibility that prevents me from trying to formally define an objective Bayesian analysis as one that has "correct" frequentist performance.

Renormalized Bayes that achieves frequentist goals sounds fun. I still am most excited, however, when the frequentist answer doesn't just use Bayesian input, but is also a Bayesian answer itself. My concern is less with the incorporation of prior information into the frequentist answer, but rather in ensuring that the frequentist answer is sound conditionally, and the only way I know to achieve this is to have the answer simultaneously be a Bayesian answer.

10 A closing comment

I do completely appreciate the role of subjective Bayesian analysis in sorting out the correct way to think about confusing situations. Much of the discussion on the subjective side of the debate has emphasized this point. Thus, in terms of philosophy of reasoning, I am a complete subjective Bayesian. But my view of statistics is that it is much broader than simply a philosophy of reasoning; it is also a vast collection of amazingly effective techniques for dealing with data, techniques that can not be easily pigeon-holed. (Of course, there is also a large collection of techniques we need to discard, and Bayesian thinking works wonders in revealing the bad ones.)

As Bayesian analysis relates to this vastness of statistics, I believe that foundations have little to do with the subjective/objective debate – i.e., to the type of prior distributions that are used. I put Bayesian analysis at the core of my foundations of the field of statistics because (i) it conditions properly on the data, and (ii) it correctly processes the uncertainties arising from the wide variety of inputs that may go into a statistical analysis, both of which are highly problematic from other approaches to statistics. The frequentist principle (in repeated use of a particular procedure, one wants the long-run behavior to be good) is also part of my foundations, but is less central because it does not handle (i) or (ii) particularly well.

References

Bayarri, M. J., Berger, J. O., Higdon, D., Kennedy, M. C., Kottas, A., Paulo, R., Sacks, J., Cafeo, J. A., Cavendish, J., Lin, C. H., and Tu, J. (2002). "A Framework for Validation of Computer Models." Technical Report 128, National Institute of Statistical Sciences, http://www.niss.org/technicalreports/tr128.pdf.

Bayarri, M. J., Berger, J. O., Kennedy, M., Kottas, A., Paulo, R., Sacks, J., Cafeo, J. A., Lin, C. H., and Tu, J. (2005b). "Validation of a Computer Model for Vehicle Collision." Technical report, National Institute of Statistical Sciences.

Bayarri, M. J., Berger, J. O., Paulo, R., Sacks, J., Cafeo, J. A., Cavendish, J., Lin, C. H., and Tu, J. (2005a). "A Framework for Validation of Computer Models." Technical report, National Institute of Statistical Sciences.

Gleser, L. and Hwang, J. (1987). "The nonexistence of $100(1-\alpha)\%$ confidence sets of finite expected diameter in errors-in-variable and related models." *Ann. Statist.*, 15: 1351–1362.