Bayesian Analysis (2015) **10**, Number 1, pp. 243–246

Rejoinder*

James O. Berger†, Jose M. Bernardo‡, and Dongchu Sun§

Our thanks to all the discussants for their insightful observations and comments. We respond to their discussions in turn.

1 Response to Datta and Liseo

We agree that Method 3 is preferable to Method 2, in that it is not dependent on the specification of a collection of quantities of interest and, hence, need only be determined once (and not separately for each potential user of a model). It is because a hierarchical embedding is not always available that we introduce the other methods as possible solutions.

We found the discussion of the multinomial example interesting, with numerous additional insights being provided. Likewise the additional material on the geometric averaging approach was enlightening, especially the nice lemma showing that, if a collection of priors all yield proper posteriors, then their geometric average also yields a proper posterior. This certainly strengthens the argument that geometric averaging is superior to arithmetic averaging in the search for an overall prior.

The moral of the amusing anecdote is indeed sound, and can be attempted to be implemented even when there is no hierarchical embedding available. For instance, Berger and Sun (2008) considered 21 different derived parameters for the five-parameter bivariate normal distribution, seeking a prior that was good 'on average' for the 21 parameters.

2 Response to Mendoza and Gutiérrez-Peña

The discussants highlight the importance of cataloguing those situations in which there is a common reference prior for all the parameters of a model and give useful references that could be a starting point for identifying additional such situations. But they then, interestingly, question whether this is sufficient, especially when the number, m, of quantities of interest exceeds the number, k, of parameters in the model.

Section 3.1 highlights one such situation: there is a common reference prior for μ and σ from the normal model but this cannot necessarily be claimed to be the overall objective prior because the reference prior for μ/σ is different. This simple example suggests that one can probably never have an overall objective prior that is optimal for everything and that just having it be reasonable for everything (of interest) might be the

*Main article DOI: 10.1214/14-BA915.
†Duke University, USA and King Abdulaziz University, Saudi Arabia, berger@stat.duke.edu
‡Universitat de València, Spain, jose.m.bernardo@uv.es
§University of Missouri-Columbia, USA, sund@missouri.edu

 DOI: 10.1214/15-BA943