Quantifying Observed Prior Impact

David Jones, Texas A&M University

It is often important to understand the impact of prior information, especially when multiple models or analyses from multiple instruments are being compared or combined. We distinguish two questions: (i) how much information does the prior contain, and (ii) what is the effect of the prior. Several measures have been proposed for quantifying effective prior sample size, for example, Clarke (1996) and Morita et al. (2008). However, these measures typically ignore the likelihood for the inference currently at hand, and therefore address (i) rather than (ii). Since in practice (ii) is of great concern, Reimherr et al. (2014) introduced a new class of effective prior sample size measures based on prior-likelihood discordance. We take this idea further towards its natural Bayesian conclusion by proposing measures of effective prior sample size that not only incorporate the general mathematical form of the likelihood but also the specific data at hand. Thus, our measures do not average across datasets from the working model, but condition on the current observed data. Consequently, our measures can be highly variable, but we demonstrate that this is because the impact of a prior can intrinsically be highly variable. Our measures are Bayes estimates of meaningful quantities and well communicate the extent to which inference is determined by the prior, or framed differently, the amount of effort saved due to having prior information.