Table 1. Frequentist Bayes factor and the Bayes factors under priors (a)-(c)for dyestuff data.

Frequentist	(a)	(b)	(c)
6.568	4.92424	8.67469	3.02869

Note that all the three Bayes factors constructed using noninformative priors (a)-(c) and the frequentist Bayes factor is a function of b. Figures 1 and 2 plot logarithm of Bayes factors against b for $m = 6$ and $m = 20$ (in each case $n_0 = 5$). It is clear that there is very good reconciliation of the Bayes factors under noninformative priors (a)-(c).

ADDITIONAL REFERENCES

Box, G. E. P. and Tiao, G. C. (1973). *Bayesian Inference in Statistical Analysis.* Wiley, New York.

Chernoff, H. (1954). On the distribution of the likelihood ratio, *J. Amer. Statist. Asso.* **25**, 573-578.

Datta, G.S. and Lahiri, P. (2000). A comparison of the Bayes factors and frequentist Bayes factor for a balanced one-way random effects model, *unpublished manuscript.*

REJOINDER

Bradley Efron and Alan Gous

This article was written under the following rule of thumb: no method that's been heavily used in serious statistical practice can be entirely wrong. The rule certainly applies to Fisherian hypothesis testing, but it also applies to Jeffreys and the BIC, leaving us to worry about Figure 1. The two scales of evidence seem to be giving radically different answers, even for sample sizes as small as $n = 100$.

Our paper localizes the disagreement to coherency, in this case sample size coherency, the key distinguishing feature of modern Bayesian philosophy. The BIC, along with any other methodology that acts coherently across different sample sizes, must share Figure 1's behavior, treating the smaller hypothesis M_o ever more favorably as n increases. Fisher's theory, which is usually presented with the sample size fixed, eschews sample size coherency in favor of a more aggressive demand for statistical power.