

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Inherent Difficulties of Non-Bayesian Likelihood-based Inference, as Revealed by an Examination of a Recent book by Aitkin*

A. GELMAN¹, C.P. ROBERT^{2,3,4}, AND J. ROUSSEAU^{2,4,5}

¹Depts. of Statistics and of Political Science, Columbia University

²Université Paris-Dauphine, CEREMADE

³Institut Universitaire de France, ⁴CREST, and ⁵ENSAE

November 6, 2011

Abstract

For many decades, statisticians have made attempts to prepare the Bayesian omelette without breaking the Bayesian eggs; that is, to obtain probabilistic likelihood-based inferences without relying on informative prior distributions. A recent example is Murray Aitkin's recent book, *Statistical Inference*, which presents an approach to statistical hypothesis testing based on comparisons of posterior distributions of likelihoods under competing models. Aitkin develops and illustrates his method using some simple examples of inference from iid data and two-way tests of independence. We analyze in this note some consequences of the inferential paradigm adopted therein, discussing why the approach is incompatible with a Bayesian perspective and why we do not find it relevant for applied work.

Keywords: Foundations, likelihood, Bayesian, Bayes factor, model choice, testing of hypotheses, improper priors, coherence.

1 Introduction

For many decades, statisticians have made attempts to prepare the Bayesian omelette without breaking the Bayesian eggs; that is, to obtain probabilistic likelihood-based inferences without relying on informative prior distributions. A recent example is Murray Aitkin's recent book, *Statistical Inference*, which is the culmination of a long research program on the topic of integrated evidence, exemplified by the discussion paper of Aitkin (1991). The book, subtitled *An*

*gelman@stat.columbia.edu, xian@ceremade.dauphine.fr, rousseau@ensae.fr

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Integrated Bayesian/Likelihood Approach, proposes handling statistical hypothesis testing and model selection via comparisons of posterior distributions of likelihood functions under the competing models or via the posterior distribution of the likelihood ratios corresponding to those models. (The essence of the proposal is detailed in Section 2.) Instead of comparing Bayes factors or performing posterior predictive checks (comparing observed data to posterior replicated pseudo-datasets), *Statistical Inference* recommends a fusion between likelihood and Bayesian paradigms that allows for the perpetuation of noninformative priors in testing settings where standard Bayesian practice prohibits their usage (DeGroot, 1973) or requires an extended decision-theoretic framework (Bernardo, 2011). While we appreciate the considerable effort made by Aitkin to place his theory within a Bayesian framework, we remain unconvinced of the said coherence, for reasons exposed in this note.




Figure 1: *Cover of Statistical Inference*

From our Bayesian perspective, and for several distinct reasons detailed in the present note, integrated Bayesian/likelihood inference cannot fit within the philosophy of Bayesian inference. Aitkin’s commendable attempt at creating a framework that incorporate the use of arbitrary noninformative priors in model choice procedures is thus incoherent in this Bayesian respect. When using improper priors lead to meaningless Bayesian procedures for posterior model comparison, we see this as a sign that the Bayesian model will not work for the problem at hand. Rather than trying at all cost to keep the offending model and define marginal posterior probabilities by fiat (whether by BIC, DIC, intrinsic Bayes factors, or posterior likelihoods), we prefer to follow the full logic of Bayesian inference and recognize that, when one’s Bayesian approach leads to a dead end, one must change either one’s methodologies or one’s beliefs (or both). Bayesians, both subjective and objective, have long recognized the need for tuning, expanding, or otherwise altering a model in light of its predictions (see, for example, Good, 1950 and Jaynes, 2003), and we view undefined Bayes factors as an example where otherwise useful methods are being extended beyond their applicability. To try to work around such problems without altering the prior distribution is, we believe, an abandonment of Bayesian principles and, more importantly, an abandoned opportunity for model improvement.

The criticisms found in the current review are therefore not limited to Aitkin’s book; they also apply to previous patches such as the deviance information criterion (DIC) of Spiegelhalter et al. (2002) (which also uses a “posterior” expectation of the log-likelihood) and the pseudo-posteriors of Geisser and Eddy (1979)(which make an extensive use of the data in their product of predictives).

Unlike the author, who has felt the call to construct a partly new Aitkin, 1991 if tentatively unifying foundation for statistical inference, we have the luxury of feeling that we already live in a comfortable (even if not flawless) inferential house. Thus, we come to Aitkin’s book not with a perceived need to rebuild but rather with a view toward strengthening the potential shakiness

1
2
3
4
5
6
7
8
9 of the pillars that support our own inferences. A key question when looking
10 at any method for probabilistic inference that is not fully Bayesian is: For
11 the applied problems that interest us, does the proposed new approach achieve
12 better performances than our existing methods? Our answer, to which we arrive
13 after careful thought, is no.

14 As an evaluation of the ideas found in *Statistical Inference*, the criticisms
15 found in this review are inherently limited. We do not claim here that Aitkin's
16 approach is wrong *per se* merely that it does not fit within our inferential
17 methodology, namely Bayesian statistics, despite using Bayesian tools. We ac-
18 knowledge that statistical methods do not, and most likely never will, form
19 a seamless logical structure. It may thus very well be that the approach of
20 comparing posterior distributions of likelihoods could be useful for some ac-
21 tual applications, and perhaps Aitkin's book will inspire future researchers to
22 demonstrate this.

23 *Statistical Inference* begins with a crisp review of frequentist, likelihood and
24 Bayesian approaches to inference and then proceeds to the main issue: intro-
25 ducing the "integrated Bayes/likelihood approach", first described in Chapter 2.
26 Much of the remaining methodological material appears in Chapters 4 ("Unified
27 analysis of finite populations") and 7 ("Goodness of fit and model diagnostics").
28 The remaining chapters apply Aitkin's principles to various examples. The
29 present article discusses the basic ideas in *Statistical Inference*, then consider
30 the relevance of Aitkin's methodology within the Bayesian paradigm.
31
32

33 2 A small change in the paradigm

34 2.1 Posterior likelihood

35
36
37 *"This quite small change to standard Bayesian analysis allows a very gen-
38 eral approach to a wide range of apparently different inference problems;
39 a particular advantage of the approach is that it can use the same nonin-
40 formative priors."* *Statistical Inference*, page *xiii*

41
42 The "quite small change" advocated by *Statistical Inference* consists in en-
43 visioning the likelihood function $L(\theta, x)$ as a generic function of the parameter
44 θ that can be processed a posteriori (that is, with a distribution induced by the
45 posterior $\pi(\theta|x)$), hence allowing for (posterior) cdf, mean, variance and quan-
46 tiles. In particular, the central tool for Aitkin's model fit is the "posterior cdf"
47 of the likelihood,

$$48 F(z) = \Pr^\pi(L(\theta, x) > z|x).$$

49
50 As argued by the author (Chapter 2, page 21), this "small change" in perspective
51 has several appealing features:

- 52
53 – The approach is general and allows to resolve the difficulties with the
54 Bayesian processing of point null hypotheses, being defined solely by the
55 Bayesian model associated with $L(\theta, x)$;

- The approach allows for the use of generic noninformative and improper priors, again by being relative to a single model;
- The approach handles more naturally the “vexed question of model fit”, still for the same reason;
- The approach is “simple.”

As noted above, the setting is quite similar to Spiegelhalter et al.’s (2002) DIC in that the deviance $D(\theta) = -2\log(f(x|\theta))$ is a renaming of the likelihood and is considered “a posteriori” both in $\bar{D} = \mathbb{E}[D(\theta)|x]$ and in $p_D = \bar{D} - D(\hat{\theta})$, where $\hat{\theta}$ is a Bayesian estimator of θ , since

$$DIC = p_D + \bar{D}.$$

The discussion of Spiegelhalter et al. (2002) made this point clear, see in particular Dawid (2002), even though the authors disagreed. Plummer (2008) make a similarly ambiguous proposal that also relates to Geisser and Eddy (1979) by its usage of cross-validation quantities.

We however dispute both the appropriateness and the magnitude of the change advocated in *Statistical Inference* and show below why, in our opinion, this shift in paradigm constitutes a new branch of statistical inference, differing from Bayesian analysis on many points. First, using priors and posteriors is no guarantee that inference is Bayesian (Seidenfeld, 1992). Empirical Bayes techniques are witnesses of this (Robbins, 1964, Carlin and Louis, 2008). Aitkin’s key departure from Bayesian principles means that his procedure has to be validated on its own, rather than benefiting from the coherence inherent to Bayesian procedures. The practical advantage of the likelihood/Bayesian approach may be convenience, but the drawback is that the method pushes both the user and the statistician *away* from progress in model building.¹

We envision Bayesian data analysis as comprising three steps: (1) model building, (2) inference, and (3) model checking. In particular, we view steps (2) and (3) as separate. Inference works well, with many exciting developments still in the coming, handling complex models, leading to an unlimited range of applications, and a partial integration with classical approaches (as in the empirical Bayes work of Efron and Morris, 1975, or more recently the similarities between hierarchical Bayes and frequentist false discovery rates discussed by

¹One might argue that, in practice, almost all Bayesians are subject to our criticism of “using models that make nonsensical predictions.” For example, Gelman et al. (2003) and Marin and Robert (2007) are full of noninformative priors. Our criticism here, though, is not of noninformative priors in general but of incoherent predictions about *quantities of interest*. In particular, noninformative priors can often (but not always!) give reasonable inferences about parameters θ within a model, even while giving meaningless (or at least not universally accepted) values for marginal likelihoods that are needed for Bayesian model comparison. It does when interest shifts from $\Pr(\theta|x, H)$ to $\Pr(H|x)$ that the Bayesian must set aside most of noninformative $\pi(\theta|H)$ and, perhaps reluctantly, set up an informative model. See, e.g., Liang et al. (2008) and Johnson and Rossell (2010) for some current perspectives on Bayesian model choice using noninformative priors.

1
2
3
4
5
6
7
8
9 Efron, 2010), causal inference, machine learning, and other aims and methods
10 of statistical inference.

11 Even in the face of all this progress on inference, Bayesian model checking
12 remains a bit of an anomaly, with the three leading Bayesian approaches be-
13 ing Bayes factors, posterior predictive checks, and comparisons of models based
14 on prediction error and other loss-based measures. (Decision-theoretic analy-
15 ses as in (Bernardo, 2011), while intellectually convincing, have not gained the
16 same amount of popularity.) Unfortunately, as Aitkin points out, none of these
17 model checking methods works completely smoothly: Bayes factors depend on
18 aspects of a model that are untestable and are commonly assigned arbitrarily;
19 posterior predictive checks are, in general, “conservative” in the sense of pro-
20 ducing p -values whose probability distributions are concentrated near 0.5; and
21 prediction error measures (which include cross-validation and DIC) require the
22 user to divide data into test and validation sets, lest they use the data twice (a
23 point discussed immediately below). The setting is even bleaker when trying to
24 incorporate noninformative priors (Gelman et al., 2003, Robert, 2001) and new
25 proposals are clearly of interest.
26
27

28 2.2 “Using the data twice” 29

30 *“A persistent criticism of the posterior likelihood approach (...) has been*
31 *based on the claim that these approaches are ‘using the data twice,’ or are*
32 *‘violating temporal coherence.’* Statistical Inference, page 48

33 “Using the data twice” is not our main reservation about the method—if
34 only because this is a rather vague concept. Obviously, one could criticize the
35 use of the “posterior expectation” of the likelihood as being the ratio of the
36 marginal of the twice replicated data over the marginal of the original data,
37

$$38 \mathbb{E}[L(\theta, x)|x] = \int L(\theta, x)\pi(\theta|x) d\theta = \frac{m(x, x)}{m(x)},$$

39
40 similar to Aitkin (1991) (a criticism clearly expressed in the discussion therein).
41 However, a more fundamental issue is that the “posterior” distribution of the
42 likelihood function cannot be justified from a Bayesian perspective. *Statistical*
43 *Inference* stays away from decision-theory (as stated on page xiv) so there is
44 no derivation based on a loss function or such. Our primary difficulty with the
45 integrated likelihood idea (and DIC as well) is (a) that the likelihood function
46 does not exist a priori and (b) that it requires a joint distribution to be properly
47 defined in the case of model comparison. The case for (a) is arguable, as Aitkin
48 would presumably contest that there exists a joint distribution on the likelihood,
49 even though the case of an improper prior stands out (see below). We still see
50 the concept of a posterior probability that the likelihood ratio is larger than
51 1 as meaningless. The case for (b) is more clear-cut in that when considering
52 two models, hence a likelihood ratio, a Bayesian analysis does require a joint
53 distribution on the two sets of parameters to reach a decision, even though in
54 the end only one set will be used. As detailed below in Section 4, this point is
55
56
57
58

1
2
3
4
5
6
7
8
9 related with the introduction of pseudo-priors by Carlin and Chib (1995) who
10 needed arbitrary defined prior distributions on the parameters that do not exist.

11 In the specific case of an improper prior, Aitkin's approach cannot be val-
12 idated in a probability setting for the reason that there is no joint probability
13 on (θ, x) . Obviously, one could always advance that the whole issue is irrele-
14 vant since improper priors do not stand within probability theory. However,
15 improper priors do stand within the Bayesian framework, as demonstrated for
16 instance by Hartigan (1983) and it is easy to give those priors an exact meaning.
17 When the data are made of n iid observations $x^n = (x_1, \dots, x_n)$ from f_θ and
18 an improper prior π is used on θ , we can consider a *training sample* (Smith and
19 Spiegelhalter, 1982) $x^{(l)}$, with $(l) \subset \{1, \dots, n\}$ such that
20

$$21 \int f(x^{(l)}|\theta) d\pi(\theta) < \infty \quad (l \leq n).$$

22
23
24 If we construct a probability distribution on θ by

$$25 \pi_{x^{(l)}}(\theta) \propto \pi(\theta)f(x^{(l)}|\theta),$$

26
27
28 the posterior distribution associated with this distribution and the remainder
29 of the sample $x^{(-l)}$ is given by

$$30 \pi_{x^{(l)}}(\theta|x^{(-l)}) \propto \pi(\theta)f(x^n|\theta), \quad x^{(-l)} = \{x_i, i \notin (l)\}.$$

31
32
33 This distribution is independent from the choice of the training sample; it only
34 depends on the likelihood of the whole data x^n and it therefore leads to a
35 non-ambiguous posterior distribution² on θ . However, as is well known, this
36 construction does not lead to produce a joint distribution on (x^n, θ) , which
37 would be required to give a meaning to Aitkin's integrated likelihood. Therefore,
38 his approach cannot cover the case of improper priors within a probabilistic
39 framework and thus fails to solve the very difficulty with noninformative priors
40 it aimed at solving. This is further illustrated by the use of Haldane's prior in
41 Chapter 4 of *Statistical Inference*, despite it not allowing for empty cells in a
42 contingency table (Jeffreys, 1939).
43
44

45 3 Posterior probability on the posterior proba- 46 bilities 47

48
49 *"The p-value is equal to the posterior probability that the likelihood ra-
50 tio, for null hypothesis to alternative, is greater than 1 (...)*The posterior
51 *probability is p that the posterior probability of H_0 is greater than 0.5."*
52 *Statistical Inference, pages 42-43*

53 ²Obvious extensions to the case of independent but non iid data or of exchangeable data
54 lead to the same interpretation. The case of dependent data is more delicate, but similar
55 interpretation can still be considered.
56
57
58

1
2
3
4
5
6
7
8
9
10 Those two equivalent statements show that it is difficult to give a Bayesian
11 interpretation to Aitkin’s method, since the two “posterior probabilities” quoted
12 above are incompatible. Indeed, a fundamental Bayesian property is that the
13 posterior probability of an event related with the parameters of the model is
14 not a random quantity but a number. To consider the “posterior probability
15 of the posterior probability” means we are exiting the Bayesian domain, both
16 from logical and philosophical viewpoints.

17 In Chapter 2, Aitkin exposes his (foundational) reasons for choosing this
18 new approach by integrated Bayes/likelihood. His criticism of Bayes factors is
19 based on several points we feel useful to reproduce here:
20

- 21 (i). “Have we really eliminated the uncertainty about the model parameters
22 by integration? The integrated likelihood (...) is the expected value of the
23 likelihood. But what of the prior variance of the likelihood?” (page 47).
24
- 25 (ii). “Any expectation with respect to the prior implies that the data has not
26 yet been observed (...) So the “integrated likelihood” is the joint distribu-
27 tion of random variables drawn by a two-stage process. (...) The marginal
28 distribution of these random variables is not the same as the distribution
29 of Y (...) and does not bear on the question of the value of θ in that
30 population” (page 47).
31
- 32 (iii). “We cannot use an improper prior to compute the integrated likelihood.
33 This eliminate the usual improper noninformative priors widely used in
34 posterior inference.” (page 47).
35
- 36 (iv). “Any parameters in the priors (...) will affect the value of the integrated
37 likelihood and this effect does not disappear with increasing sample size”
38 (page 47).
39
- 40 (v). “The Bayes factor is equal to the posterior mean of the likelihood ratio
41 between the models” [*meaning under the full model posterior*] (page 48).
42
- 43 (vi). “The Bayes factor diverges as the prior becomes diffuse. (...) This prop-
44 erty of the Bayes factor has been known since the Lindley/Bartlett paradox
45 of 1957” (page 48).

46 The representation (i) of the “integrated” (or marginal) likelihood as an
47 expectation under the prior
48

49
50
$$m(x) = \int L(\theta, x)\pi(\theta) d\theta = \mathbb{E}^\pi[L(\theta, x)]$$

51

52 is unassailable and is for instance used as a starting point for motivating the
53 nested sampling method (Skilling, 2006, Chopin and Robert, 2010). This does
54 not imply that the extension to the variance or to any other moment stated
55 in (i) has a similar meaning, nor that the move to the expectation under the
56 posterior is valid within the Bayesian paradigm. While the difficulty (iii) with
57
58

improper priors is real, and while the impact of the prior modelling (iv) may have a lingering effect, the other points can be easily rejected on the ground that the posterior distribution of the likelihood is meaningless within a Bayesian perspective. This criticism is anticipated by Aitkin who protests on pages 48-49 that, given point (v), the posterior distribution must be “meaningful,” since the posterior mean is “meaningful”, but the interpretation of the Bayes factor as a “posterior mean” is only an interpretation of an existing integral (in the specific case of nested models), it does not give any validation to the analysis. (The marginal likelihood may similarly be interpreted as a prior mean, despite depending on the observation x , as in the nested sampling perspective. More generally, bridge sampling techniques also exploit those multiple representations of a ratio of integrals, Gelman and Meng, 1998.) One could just as well take (ii) above as an argument *against* the integrated likelihood/Bayes perspective.

4 Products of posteriors

In the case of unrelated models to be compared, the fundamental theoretical argument against using posterior distributions of the likelihoods and of related terms is that the approach leads to parallel and separate simulations from the posteriors under each model. *Statistical Inference* recommends that models be compared via the distribution of the likelihood ratio values,

$$L_i(\theta_i|x) / L_k(\theta_k|x),$$

where the θ_i 's and θ_k 's are drawn from the respective posteriors. This choice is similar to Scott's (2002) and to Congdon's (2006) mistaken solutions exposed in Robert and Marin (2008), in that MCMC simulations are run for each model separately and the resulting samples are then gathered together to produce either the posterior expectation (in Scott's, 2002, case) or the posterior distribution (for the current paper) of

$$\rho_i L(\theta_i|x) / \sum_k \rho_k L(\theta_k|x),$$

which do not correspond to genuine Bayesian solutions (see Robert and Marin, 2008). Again, this is not as much because the dataset x is used repeatedly in this process (since reversible MCMC produces as well separate samples from the different posteriors) as the fundamental lack of a common joint distribution that is needed in the Bayesian framework. This means, e.g., that the integrated likelihood/Bayes technology is producing samples from the product of the posteriors (a product that clearly is not defined in a Bayesian framework) instead of using pseudo-priors as in Carlin and Chib (1995), i.e. of considering a joint posterior on (θ_1, θ_2) , which is [proportional to]

$$p_1 m_1(x) \pi_1(\theta_1|x) \pi_2(\theta_2) + p_2 m_2(x) \pi_2(\theta_2|x) \pi_1(\theta_1). \quad (1)$$

This makes a difference in the outcome, as illustrated in Figure 2, which compares the distribution of the likelihood ratio under the true posterior and under the product of posteriors, when assessing the fit of a Poisson model against the fit of a binomial model with $m = 5$ trials, for the observation $x = 3$. The joint simulation produces a much more supportive argument in favor of the binomial model, when compared with the product of the posteriors. (Again, this is inherently the flaw found in the reasoning leading to Scott's, 2002, and Congdon's, 2006, methods for approximating Bayes factors.)

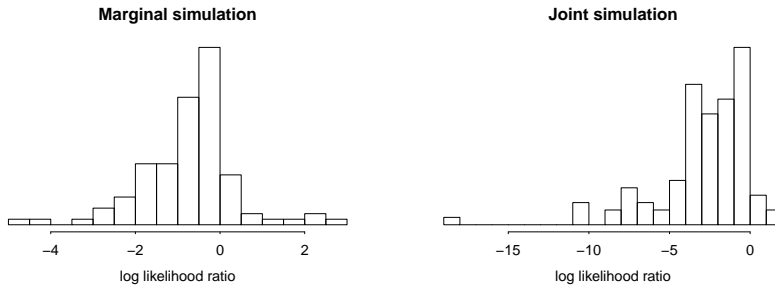


Figure 2: Comparison of the distribution of the likelihood ratio under the correct joint posterior and under the product of the model-based posteriors, when assessing a Poisson model against a binomial with $m = 5$ trials, for $x = 3$. The joint simulation produces a much more supportive argument in favor of the negative binomial model, when compared with the product of the posteriors.

Although we do not advocate its use, a Bayesian version of Aitkin's proposal can be constructed based on the following loss function that evaluates the estimation of the model index j based on the values of the parameters under both models and on the observation x :

$$L(\delta, (j, \theta_j, \theta_{-j})) = \mathbb{I}_{\delta=1} \mathbb{I}_{f_2(x|\theta_2) > f_1(x|\theta_1)} + \mathbb{I}_{\delta=2} \mathbf{I}_{f_2(x|\theta_2) < f_1(x|\theta_1)}. \quad (2)$$

Here $\delta = j$ means that model j is chosen, and $f_j(\cdot|\theta_j)$ denotes the likelihood under model j . Under this loss, the Bayes (optimal) solution is

$$\delta^\pi(x) = \begin{cases} 1 & \text{if } \Pr^\pi [f_2(x|\theta_2) < f_1(x|\theta_1)|x] > 1/2 \\ 2 & \text{otherwise,} \end{cases}$$

which depends on the joint posterior distribution (1) on (θ_1, θ_2) , thus differs from Aitkin's solution. We have

$$\begin{aligned} \Pr^\pi [f_2(x|\theta_2) < f_1(x|\theta_1)|x] &= \pi(\mathcal{M}_1|x) \int_{\Theta_2} \Pr^{\pi_1} [l^1(\theta_1) > l^2(\theta_2)|x, \theta_2] d\pi_2(\theta_2) \\ &\quad + \pi(\mathcal{M}_2|x) \int_{\Theta_1} \Pr^{\pi_2} [l^1(\theta_1) > l^2(\theta_2)|x, \theta_1] d\pi_1(\theta_1), \end{aligned}$$

where l^1 and l^2 denote the log-likelihoods and where the probabilities within the integrals are computed under $\pi_1(\theta_1|x)$ and $\pi_2(\theta_2|x)$, respectively. (Pseudo-priors as in Carlin and Chib, 1995 could be used instead of the true priors, a requirement when at least one of those priors is improper.)

An asymptotic evaluation of the above procedure is possible: consider a sample of size n , x^n . If \mathcal{M}_1 is the “true” model, then $\pi(\mathcal{M}_1|x^n) = 1 + o_p(1)$ and we have

$$\begin{aligned} \Pr^{\pi_1} [l^1(\theta_1) > l^2(\theta_2)|x^n, \theta_2] &= \Pr \left[-\mathcal{X}_{p_1}^2 > l^2(\theta_2) - l^2(\hat{\theta}_1) \right] + O_p(1/\sqrt{n}) \\ &= F_{p_1} \left[l^1(\hat{\theta}_1) - l^2(\theta_2) \right] + O_p(1/\sqrt{n}), \end{aligned}$$

with obvious notations for the corresponding log-likelihoods, p_1 the dimension of Θ_1 , $\hat{\theta}_1$ the maximum likelihood estimator of θ_1 , and $\mathcal{X}_{p_1}^2$ a chi-square random variable with p_1 degrees of freedom. Note also that, since $l^2(\theta_2) \leq l^2(\hat{\theta}_2)$,

$$l^1(\hat{\theta}_1) - l^2(\theta_2) \geq n\text{KL}(f_0, f_{\theta_2^*}) + O_p(\sqrt{n}),$$

where $\text{KL}(f, g)$ denotes the Kullback–Leibler divergence and θ_2^* denotes the *projection* of the true model on \mathcal{M}_2 : $\theta_2^* = \text{argmin}_{\theta_2} \text{KL}(f_0, f_{\theta_2})$, we have

$$\Pr^\pi [f(x^n|\theta_2) < f(x^n|\theta_1)|x^n] = 1 + o_p(1).$$

By symmetry, the same asymptotic consistency occurs under model \mathcal{M}_2 . On the opposite, Aitkin’s approach leads (at least in regular models) to the approximation

$$\Pr[\mathcal{X}_{p_2}^2 - \mathcal{X}_{p_1}^2 > l^2(\hat{\theta}_2) - l^1(\hat{\theta}_1)],$$

where the $\mathcal{X}_{p_2}^2$ and $\mathcal{X}_{p_1}^2$ random variables are independent, hence producing quite a different result that depends on the asymptotic behavior of the likelihood ratio. Note that for both approaches to be equivalent one would need a pseudo-prior for \mathcal{M}_2 (resp. \mathcal{M}_1 if \mathcal{M}_2 were *true*) as tight around the maximum likelihood as the posterior $\pi_2(\theta_2|x^n)$, which would be equivalent to some kind of empirical Bayes type of procedure.

Furthermore, in the case of embedded models, \mathcal{M}_2 and $\mathcal{M}_1 \subset \mathcal{M}_2$, Aitkin’s approach can be given a probabilistic interpretation. To this effect, we write the parameter under \mathcal{M}_1 as (θ_1, ψ_0) , ψ_0 being a fixed known quantity, and under \mathcal{M}_2 as $\theta_2 = (\theta_1, \psi)$, so that comparing \mathcal{M}_1 with \mathcal{M}_2 corresponds to testing the null hypothesis $\psi = \psi_0$. Aitkin does not impose a positive prior probability on \mathcal{M}_1 , since his prior only bears on \mathcal{M}_2 (in a spirit close to the Savage-Dickey representation, see Marin and Robert, 2010). His approach is therefore similar to the inversion of a confidence region into a testing procedure (or vice-versa). Under the model $\mathcal{M}_1 \subset \mathcal{M}_2$, denoting by $l(\theta, \psi)$ the log-likelihood of the bigger model,

$$\begin{aligned} \Pr^\pi [l(\theta_1, \psi_0) > l(\theta_1, \psi)|x^n] &\approx \Pr \left[\mathcal{X}_{p_2-p_1}^2 > -l(\hat{\theta}_1(\psi_0), \psi_0) + l(\hat{\theta}_1, \hat{\psi}) \right] \\ &\approx 1 - F_{p_2-p_1} [-l(\hat{\theta}_1(\psi_0), \psi_0) + l(\hat{\theta}_1, \hat{\psi})], \end{aligned}$$

1
2
3
4
5
6
7
8
9 which is the approximate p -value associated with the likelihood ratio test. There-
10 fore, the aim of this approach seems to be, at least for embedded models where
11 the Bernstein–von Mises theorem holds for the posterior distribution, to con-
12 struct a *Bayesian* procedure reproducing the p -value associated with the likeli-
13 hood ratio test. From a frequentist point of view it is of interest to see that the
14 posterior probability of the likelihood ratio being greater than one is approxi-
15 mately a p -value, at least in cases when the Bernstein-von Mises theorem holds,
16 e.g. for embedded models and proper priors. This p -value can then be given
17 a finite-sample meaning (under the above restrictions), however it seems more
18 interesting from a frequentist perspective than from a Bayesian one.³ From a
19 Bayesian decision-theoretic viewpoint, this is even more dubious, since the loss
20 function (2) is difficult to interpret and to justify.

21
22 *“Without a specific alternative, the best we can do is to make posterior*
23 *probability statements about μ and transfer these to the posterior distribu-*
24 *tion of the likelihood ratio (..) There cannot be strong evidence in favor*
25 *of a point null hypothesis against a general alternative hypothesis.”* Sta-
26 *tistical Inference, pages 42–44*

27 We further note that, once *Statistical Inference* has set the principle of using
28 the posterior distribution of the likelihood ratio (or rather of the divergence
29 difference since this is at least symmetric in both hypotheses), there is a whole
30 range of outputs available including confidence intervals on the difference, for
31 checking whether or not they contain zero. From our (Bayesian) perspective,
32 this solution (a) is not Bayesian for reasons exposed above, (b) is not parame-
33 terization invariant, and (c) relies once again on an arbitrary confidence level.
34

35 36 5 Misrepresentations 37

38 We have focused in this review on Aitkin’s proposals rather than on his char-
39 acterizations of other statistical methods. In a few places, however, we believe
40 that there have been some unfortunate confusions from his part.

41 On page 22, Aitkin describes Bayesian posterior distributions as “formally a
42 measure of personal uncertainty about the model parameter,” a statement that
43 we believe holds generally only under a definition of “personal” that is so broad
44 as to be meaningless. As we have discussed elsewhere (Gelman, 2008), Bayesian
45 probabilities can be viewed as “subjective” or “personal” but this is not neces-
46 sary. Or, to put it another way, if you want to label my posterior distribution as
47 “personal” because it is based on my personal choice of prior distribution, you
48 should also label inferences from the proportional hazards model as “personal”
49 because it is based on the user’s choice of the parameterization of Cox (1972);
50 you should also label any linear regression (classical or otherwise) as “personal”
51 as based on the individual’s choice of predictors and assumptions of additivity,
52 linearity, variance function, and error distribution; and so on for all but the very
53 simplest models in existence.
54

55 ³See Chapter 7 of Gelman et al. (2003) for a fully Bayesian treatment of finite-sample
56 inference.
57

1
2
3
4
5
6
7
8
9 In a nearly century-long tradition in statistics, any probability model is
10 sharply divided into “likelihood” (which is considered to be objective and, in
11 textbook presentations, is often simply given as part of the mathematical spec-
12 ification of the problem) and “prior” (a dangerously subjective entity to which
13 the statistical researcher is encouraged to pour all of his or her pent-up skep-
14 ticism). This may be a tradition but it has no logical basis. If writers such as
15 Aitkin wish to consider their likelihoods as objective and consider their priors as
16 subjective, that is their privilege. But we would prefer them to restrain them-
17 selves when characterizing the models as others. It would be polite to either
18 tentatively accept the objectivity of others’ models or, contrariwise, to gallantly
19 affirm the subjectivity of one’s own choices.

20 Aitkin also mischaracterizes hierarchical models, writing “It is important
21 not to interpret the prior as in some sense a *model for nature* [italics in the
22 original] that nature has used a random process to draw a parameter value from
23 a higher distribution of parameter values . . .” On the contrary, that is exactly
24 how we interpret the prior distribution in the ideal case. Admittedly, we do
25 not generally approach this ideal (except in settings such as genetics where the
26 population distribution of parameters has a clear sampling distribution), just
27 as in practice the error terms in our regression models do not capture the true
28 distribution of errors. Despite these imperfections, we believe that it can often
29 be helpful to interpret the prior as a model for the parameter-generation process
30 and to improve this model where appropriate.

31 32 33 34 **6 Contributions of the book**

35
36 *Statistical Inference* points out several important facts that are individually
37 known well (but perhaps not well enough!), but by putting them all in one
38 place it foregrounds the difficulty or impossibility of putting all the different
39 approaches to model checking in one place. We all know that the p -value is in
40 no way the posterior probability of a null hypothesis being true; in addition,
41 Bayes factors as generally practiced correspond to no actual probability model.
42 Also, it is well-known that the so-called harmonic mean approach to calculating
43 Bayes factors is inherently unstable, to the extent that in the situations where
44 it does “work,” it works by implicitly integrating over a space different from
45 that of its nominal model.

46 Yes, we all know these things, but as is often the case with scientific anoma-
47 lies, they are associated with such a high level of discomfort that many re-
48 searchers tend to forget the problems or try to finesse them. It is refreshing to
49 see the anomalies laid out so clearly.

50 At some points, however, Aitkin disappoints. For example, at the end of Sec-
51 tion 7.2, he writes: “In the remaining sections of this chapter, we first consider
52 the posterior predictive p -value and point out difficulties with the posterior pre-
53 dictive distribution which closely parallel those of Bayes factors.” He follows up
54 with a section entitled “The posterior predictive distribution,” which concludes
55 with an example that he writes “should be a matter of *serious* concern [em-

1
2
3
4
5
6
7
8
9 phasis in original] to those using posterior predictive distributions for predictive
10 probability statements.”

11 What is this example of serious concern? It is an imaginary problem in which
12 he observes 1 success in 10 independent trials and then is asked to compute the
13 probability of getting at most 2 successes in 20 more trials from the same pro-
14 cess. *Statistical Inference* assumes a uniform prior distribution on the success
15 probability and yields a predictive probability of 0.447, which, to him, “looks a
16 vastly optimistic and unsound statement.” Here, we think Aitkin should take
17 Bayes a bit more seriously. If you think this predictive probability is unsound,
18 there should be some aspect of the prior distribution or the likelihood that is
19 unsound as well. This is what Good (1950) called “the device of imaginary re-
20 sults.” We suggest that, rather than abandoning highly effective methods based
21 on predictive distributions, Aitkin should look more carefully at his predictive
22 distributions and either alter his model to fit his intuitions, alter his intuitions
23 to fit his model, or do a bit of both. This is the value of inferential coherence
24 as an ideal.
25

26 27 **7 Solving non-problems**

28
29 Several of the examples in *Statistical Inference* represent solutions to problems
30 that seem to us to be artificial or conventional tasks with no clear analogy to
31 applied work.
32

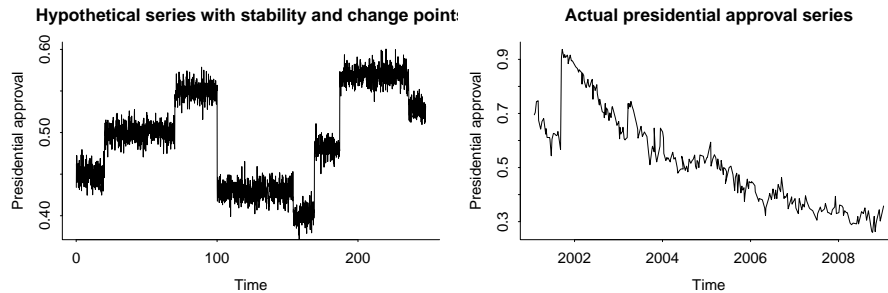
33 *“They are artificial and are expressed in terms of a survey of 100 indi-*
34 *viduals expressing support (Yes/No) for the president, before and after a*
35 *presidential address (...). The question of interest is whether there has been*
36 *a change in support between the surveys (...). We want to assess the evi-*
37 *dence for the hypothesis of equality H_1 against the alternative hypothesis*
38 *H_2 of a change.”* *Statistical Inference, page 147*

39 Based on our experience in public opinion research, this is not a real question.
40 Support for any political position is always changing. The real question is how
41 much the support has changed, or perhaps how this change is distributed across
42 the population.
43

44 A defender of Aitkin (and of classical hypothesis testing) might respond at
45 this point that, yes, everybody knows that changes are never exactly zero and
46 that we should take a more “grown-up” view of the null hypothesis, not that the
47 change is zero but that it is nearly zero. Unfortunately, the metaphorical inter-
48 pretation of hypothesis tests has problems similar to the theological doctrines
49 of the Unitarian church. Once you have abandoned literal belief in the Bible,
50 the question soon arises: why follow it at all? Similarly, once one recognizes
51 the inappropriateness of the point null hypothesis, it makes more sense not to
52 try to rehabilitate it or treat it as treasured metaphor but rather to attack our
53 statistical problems directly, in this case by performing inference on the change
54 in opinion in the population.
55

56 To be clear: we are not denying the value of hypothesis testing. In this
57 example, we find it completely reasonable to ask whether observed changes
58
59

1
2
3
4
5
6
7
8
9 are statistically significant, i.e. whether the data are consistent with a null
10 hypothesis of zero change. What we do not find reasonable is the statement
11 that “the question of interest is whether there has been a change in support.”
12



13
14
15
16
17
18
19
20
21
22
23
24
25 Figure 3: (a) Hypothetical graph of presidential approval with discrete jumps;
26 (b) presidential approval series (for George W. Bush) showing movement at
27 many different time scales. If the approval series looked like the graph on the
28 left, then Aitkin’s “question of interest” of “whether there has been a change in
29 support between the surveys” would be completely reasonable. In the context of
30 actual public opinion data, the question does not make sense; instead, we prefer
31 to think of presidential approval as a continuously-varying process.
32

33 All this is application-specific. Suppose public opinion was observed to really
34 be flat, punctuated by occasional changes, as in the left graph in Figure 3. In
35 that case, Aitkin’s question of “whether there has been a change” would be
36 well-defined and appropriate, in that we could interpret the null hypothesis of
37 no change as some minimal level of baseline variation.
38

39 Real public opinion, however, does not look like baseline noise plus jumps,
40 but rather shows continuous movement on many time scales at once, as can be
41 seen from the right graph in Figure 3, which shows actual presidential approval
42 data. In this example, we do not see Aitkin’s question as at all reasonable. Any
43 attempt to work with a null hypothesis of opinion stability will be inherently
44 arbitrary. It would make much more sense to model opinion as a continuously-
45 varying process.

46 The statistical problem here is not merely that the null hypothesis of zero
47 change is nonsensical; it is that the null is in no sense a reasonable approximation
48 to any interesting model. The sociological problem is that, from Savage (1954)
49 onward, many Bayesians have felt the need to mimic the classical null-hypothesis
50 testing framework, even where it makes no sense. Aitkin is unfortunately no
51 exception, taking a straightforward statistical question—estimating a time trend
52 in opinion—and re-expressing it as an abstracted hypothesis testing problem
53 that pulls the analyst away from any interesting political questions.
54
55
56
57
58

8 Conclusion: Why did we write this review?

“The posterior has a non-integrable spike at zero. This is equivalent to assigning zero prior probability to these unobserved values.” Statistical Inference, page 98

A skeptical (or even not so skeptical) reader might at this point ask, Why did we bother to write a detailed review of a somewhat obscure statistical method that we do not even like? Our motivation surely was not to protect the world from a dangerous idea; if anything, we suspect our review will interest some readers who otherwise would not have heard about the approach (as previously illustrated by Robert, 2010).

In 1970, a book such as *Statistical Inference* could have had a large influence in statistics. As Aitkin notes in his preface, there was a resurgence of interest in the foundations of statistics around that time, with Lindley, Dempster, Barnard, and others writing about the intersections between classical and Bayesian inference (going beyond the long-understood results of asymptotic equivalence) and researchers such as Akaike and Mallows beginning to integrate model-based and predictive approaches to inference. A glance at the influential text of Cox and Hinkley (1974) reveals that theoretical statistics at that time was focused on inference from independent data from specified sampling distributions (possibly after discarding information, as in rank-based tests), and “likelihood” was central to all these discussions.

Forty years on, a book on likelihood inference is more of a niche item. Partly this is simply part of the growth of the field—with the proliferation of books, journals, and online publications, it is much more difficult for any single book to gain prominence. More than that, though, we think statistical theory has moved away from iid analysis, toward more complex, structured problems.

That said, the foundational problems that *Statistical Inference* discusses are indeed important and they have not yet been resolved. As models get larger, the problem of “nuisance parameters” is revealed to be not a mere nuisance but rather a central fact in all methods of statistical inference. As noted above, Aitkin makes valuable points—known, but not well-enough known—about the difficulties of Bayes factors, pure likelihood, and other superficially attractive approaches to model comparison. We believe it is a natural continuation of this work to point out the problems of the integrated likelihood approach as well.

For now, we recommend model expansion, Bayes factors where reasonable, cross-validation, and predictive model checking based on graphics rather than p -values. We recognize that each of these approaches has loose ends. But, as practical idealists, we consider inferential challenges to be opportunities for model improvement with the Bayesian realm rather than motivations for a new theory of noninformative priors that takes us in uncharted territories.

References

- Aitkin, M. 1991. Posterior Bayes factors (with discussion). *J. Royal Statist. Society Series B* 53: 111–142.
- Bernardo, J. 2011. Integrated objective Bayesian estimation and hypothesis testing. In *Bayesian Statistics 9*, eds. J. Bernardo, J. Berger, M. Bayarri, D. Heckerman, A. Smith, and M. West, 1–68. Oxford: Oxford University Press.
- Carlin, B. and S. Chib. 1995. Bayesian model choice through Markov chain Monte Carlo. *J. Royal Statist. Society Series B* 57(3): 473–484.
- Carlin, B. and T. Louis. 2008. *Bayes and Empirical Bayes Methods for Data Analysis*. 3rd ed. Chapman and Hall, New York.
- Chopin, N. and C. Robert. 2010. Properties of nested sampling. *Biometrika* 97: 741–755.
- Congdon, P. 2006. Bayesian model choice based on Monte Carlo estimates of posterior model probabilities. *Comput. Stat. Data Analysis* 50: 346–357.
- Dawid, A. 2002. Discussion of “Bayesian measures of model complexity and fit” by Spiegelhalter et al. *J. Royal Statist. Society Series B* 64(2): 583–639.
- DeGroot, M. 1973. Doing what comes naturally: Interpreting a tail area as a posterior probability or as a likelihood ratio. *J. American Statist. Assoc.* 68: 966–969.
- Efron, B. 2010. The future of indirect evidence (with discussion). *Statist. Science* 25(2): 145–171.
- Efron, B. and C. Morris. 1975. Data analysis using Stein’s estimator and its generalizations. *J. American Statist. Assoc.* 70: 311–319.
- Geisser, S. and W. Eddy. 1979. A predictive approach to model selection. *J. American Statist. Assoc.* 74: 153–160.
- Gelman, A., J. Carlin, H. Stern, and D. Rubin. 2003. *Bayesian Data Analysis*. 2nd ed. New York: Chapman and Hall, New York.
- Gelman, A. and X. Meng. 1998. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statist. Science* 13: 163–185.
- Good, I. 1950. *Probability and the Weighting of Evidence*. London: Charles Griffin.
- Hartigan, J. A. 1983. *Bayes Theory*. New York: Springer-Verlag, New York.
- Jaynes, E. 2003. *Probability Theory*. Cambridge: Cambridge University Press.

- 1
2
3
4
5
6
7
8
9 Jeffreys, H. 1939. *Theory of Probability*. 1st ed. Oxford: The Clarendon Press.
- 10 Johnson, V. and D. Rossell. 2010. On the use of non-local prior densities in
11 Bayesian hypothesis tests. *J. Royal Statist. Society Series B* 72: 143–170.
- 12
13 Liang, F., R. Paulo, G. Molina, M. Clyde, and J. Berger. 2008. Mixtures of g
14 priors for Bayesian variable selection. *J. American Statist. Assoc.* 103(481):
15 410–423.
- 16
17 Marin, J. and C. Robert. 2007. *Bayesian Core*. Springer-Verlag, New York.
- 18
19 —. 2010. On resolving the Savage–Dickey paradox. *Electron. J. Statist.* 4:
20 643–654.
- 21
22 Plummer, M. 2008. Penalized loss functions for Bayesian model comparison.
23 *Biostatistics* 9(3): 523–539.
- 24
25 Robbins, H. 1964. The empirical Bayes approach to statistical decision problems.
26 *Ann. Math. Statist.* 35: 1–20.
- 27
28 Robert, C. 2001. *The Bayesian Choice*. 2nd ed. Springer-Verlag, New York.
- 29
30 —. 2010. The Search for Certainty: a critical assessment. *Bayesian Analysis*
31 5(2): 213–222. (with discussion).
- 32
33 Robert, C. and J.-M. Marin. 2008. On some difficulties with a posterior proba-
34 bility approximation technique. *Bayesian Analysis* 3(2): 427–442.
- 35
36 Savage, L. 1954. *The Foundations of Statistical Inference*. New York: John
37 Wiley.
- 38
39 Scott, S. L. 2002. Bayesian methods for hidden Markov models: recursive com-
40 puting in the 21st Century. *J. American Statist. Assoc.* 97: 337–351.
- 41
42 Seidenfeld, T. 1992. R.A. Fisher’s fiducial argument and Bayes’ theorem. *Statist.*
43 *Science* 7(3): 358–368.
- 44
45 Skilling, J. 2006. Nested sampling for general Bayesian computation. *Bayesian*
46 *Analysis* 1(4): 833–860.
- 47
48 Smith, A. and D. Spiegelhalter. 1982. Bayes factors for linear and log-linear
49 models with vague prior information. *J. Royal Statist. Society Series B* 44:
50 377–387.
- 51
52 Spiegelhalter, D. J., N. G. Best, B. P. Carlin, and A. van der Linde. 2002.
53 Bayesian measures of model complexity and fit (with discussion). *J. Royal*
54 *Statist. Society Series B* 64(2): 583–639.
- 55
56
57
58
59
60
61
62
63
64
65