

# Frequentist tests for Bayesian models

L. B. Lucy

Astrophysics Group, Blackett Laboratory, Imperial College London, Prince Consort Road, London SW7 2AZ, UK  
e-mail: l.lucy@imperial.ac.uk

Received 6 November 2015 / Accepted 1 February 2016

## ABSTRACT

Analogues of the frequentist chi-square and  $F$  tests are proposed for testing goodness-of-fit and consistency for Bayesian models. Simple examples exhibit these tests' detection of inconsistency between consecutive experiments with identical parameters, when the first experiment provides the prior for the second. In a related analysis, a quantitative measure is derived for judging the degree of tension between two different experiments with partially overlapping parameter vectors.

**Key words.** methods: data analysis – methods: statistical

## 1. Introduction

Bayesian statistical methods are now widely applied in astronomy. Of the new techniques thus introduced, model selection (or comparison) is especially notable in that it replaces the frequentist acceptance-rejection paradigm for testing hypotheses. Thus, given a data set  $D$ , there might be several hypotheses  $\{H_k\}$  that have the potential to explain  $D$ . From these, Bayesian model selection identifies the particular  $H_k$  that best explains  $D$ . The procedure is simple, though computationally demanding: starting with an arbitrary pair of the  $\{H_k\}$ , we apply the model selection machinery, discard the weaker hypothesis, replace it with one of the remaining  $H_k$ , and repeat.

This procedure usefully disposes of the weakest hypotheses, but there is no guarantee that the surviving  $H_k$  explains  $D$ . If the correct hypothesis is not included in the  $\{H_k\}$ , we are left with the “best” hypothesis but are not made aware that the search for an explanation should continue. In the context of model selection, the next step would be a comparison of this “best”  $H_k$  with the hypothesis that  $H_k$  is false. But model selection fails at this point because we cannot compute the required likelihood (Sivia & Skilling 2006, p. 84). Clearly, what is needed is a goodness-of-fit criterion for Bayesian models.

This issue is discussed by Press et al. (2007, p. 779). They note that “There are no good fully Bayesian methods for assessing goodness-of-fit ...” and go on to report that “Sensible Bayesians usually fall back to  $p$ -value tail statistics ... when they really need to know if a model is wrong”.

On the assumption that astronomers do really need to know if their models are wrong, this paper adopts a frequentist approach to testing Bayesian models. Although this approach may be immediately abhorrent to committed Bayesians, the role of the tests proposed herein is merely to provide a quantitative measure according to which Bayesians decide whether their models are satisfactory. When they are, the Bayesian inferences are presented – and with increased confidence. When not, flaws in the models or the data must be sought, with the aim of eventually achieving satisfactory Bayesian inferences.

## 2. Bayesian models

The term Bayesian model – subsequently denoted by  $\mathcal{M}$  – must now be defined. The natural definition of  $\mathcal{M}$  is that which must be specified in order that Bayesian inferences can be drawn from  $D$  – i.e., in order to compute posterior probabilities. This definition implies that, in addition to the hypothesis  $H$ , which introduces the parameter vector  $\alpha$ , the prior probability distribution  $\pi(\alpha)$  must also be included in  $\mathcal{M}$ . Thus, symbolically, we write

$$\mathcal{M} \equiv \{\pi, H\}. \quad (1)$$

It follows that different Bayesian models can share a common  $H$ . For example,  $H$  may be the hypothesis that  $D$  is due to Keplerian motion. But for the motion of a star about the Galaxy's central black hole, the appropriate  $\pi$  will differ from that for the reflex orbit of star due to a planetary companion.

A further consequence is that a failure of  $\mathcal{M}$  to explain  $D$  is not necessarily due to  $H$ : an inappropriate  $\pi$  is also a possibility.

To astronomers accustomed to working only with uniform priors, the notion that a Bayesian model's poor fit to  $D$  could be due to  $\pi$  might be surprising. A specific circumstance where  $\pi$  could be at fault arises when Bayesian methods are used to repeatedly update our knowledge of some phenomenon – e.g., the value of a fundamental constant that over the years is the subject of numerous experiments ( $X_1, \dots, X_i, \dots$ ), usually of increasing precision. With an orthodox Bayesian approach,  $X_{i+1}$  is analysed with the prior set equal to the posterior from  $X_i$ . Thus

$$\pi_{i+1}(\alpha) = p(\alpha|H, D_i). \quad (2)$$

This is the classical use of Bayes theorem to update our opinion by incorporating past experience. However, if  $X_i$  is flawed – e.g., due to an unrecognized systematic error – then this choice of  $\pi$  impacts negatively on the analysis of  $X_{i+1}$ , leading perhaps to a poor fit to  $D_{i+1}$ .

Now, since subsequent flawless experiments result in the decay of the negative impact of  $X_i$ , this recursive procedure is self-correcting, so a Bayesian might argue that the problem can be

ignored. But scientists feel obliged to resolve such discrepancies before publishing or undertaking further experiments and therefore need a statistic that quantifies any failure of  $\mathcal{M}$  to explain  $D$ .

### 3. A goodness-of-fit statistic for Bayesian models

The most widely used goodness-of-fit statistic in the frequentist approach to hypothesis testing is  $\chi^2$ , whose value is determined by the residuals between the fitted model and the data, with no input from prior knowledge. Thus,

$$\chi_0^2 = \chi^2(\alpha_0) \quad (3)$$

is the goodness-of-fit statistic for the minimum- $\chi^2$  solution  $\alpha_0$ .

A simple analogue of  $\chi_0^2$  for Bayesian models is the posterior mean of  $\chi^2(\alpha)$ ,

$$\langle \chi^2 \rangle_\pi = \int \chi^2(\alpha) p(\alpha|H, D) dV_\alpha \quad (4)$$

where the posterior distribution

$$p(\alpha|H, D) = \frac{\pi(\alpha)\mathcal{L}(\alpha|H, D)}{\int \pi(\alpha)\mathcal{L}(\alpha|H, D) dV_\alpha} \quad (5)$$

Here  $\mathcal{L}(\alpha|H, D)$  is the likelihood of  $\alpha$  given data  $D$ .

Note that since  $\langle \chi^2 \rangle_\pi$  depends on both constituents of  $\mathcal{M}$ , namely  $H$  and  $\pi$ , it has the potential to detect a poor fit due to either or both being at fault, as required by Sect. 2.

In Eq. (4) the subscript  $\pi$  is attached to  $\langle \chi^2 \rangle$  to stress that a non-trivial, informative prior is included in  $\mathcal{M}$ . On the other hand, when a uniform prior is assumed,  $\langle \chi^2 \rangle$  is independent of the prior and is then denoted by  $\langle \chi^2 \rangle_u$ .

The Bayesian goodness-of-fit statistic  $\langle \chi^2 \rangle_u$  is used in Lucy (2014, L14) to illustrate the detection of a weak second orbit in simulations of *Gaia* scans of an astrometric binary. In that case,  $H$  states that the scan residuals are due to the reflex Keplerian orbit caused by *one* invisible companion. With increasing amplitude of the second orbit, the point comes when the investigator will surely abandon  $H$  – i.e., abandon the assumption of just one companion – see Fig. 12, L14.

#### 3.1. $p$ -values

With the classical frequentist acceptance-rejection paradigm, a null hypothesis  $H_0$  is rejected on the basis of a  $p$ -value tail statistic. Thus, with the  $\chi^2$  test,  $H_0$  is rejected if  $\chi_0^2 > \chi_{\nu, \beta}^2$ , where  $p(\chi_\nu^2 > \chi_{\nu, \beta}^2) = \beta$ , and accepted otherwise. Here  $\nu = n - k$  is the number of degrees of freedom, where  $n$  is the number of measurements and  $k$  is the number of parameters introduced by  $H_0$ , and  $\beta$  is the designated  $p$ -threshold chosen by the investigator.

For a Bayesian model, a  $p$ -value can be computed from the  $\langle \chi^2 \rangle_\pi$  statistic, whose approximate distribution is derived below in Sect. 5.1. However, a sharp transition from acceptance to rejection of the null hypothesis at some designated  $p$ -value is not recommended. First,  $p$ -values overstate the strength of the evidence against  $H_0$  (e.g., Sellke et al. 2001). In particular, the value  $p = 0.05$  recommended in elementary texts does not imply strong evidence against  $H_0$ . Second, the  $p$ -value is best regarded (Sivia & Skilling 2006, p. 85) as serving a qualitative purpose, with a small value prompting us to think about alternative hypotheses. Thus, if  $\langle \chi^2 \rangle_\pi$  exceeds the chosen threshold  $\chi_{\nu, \beta}^2$ , this is a warning that something is amiss and should be investigated, with the degree of concern increasing as  $\beta$  decreases. If

the  $\beta = 0.001$  threshold is exceeded, then the investigator would be well-advised to suspect that  $\mathcal{M}$  or  $D$  is at fault.

Although statistics texts emphasize tests of  $H$  not  $D$ , astronomers know that  $D$  can be corrupted by biases or calibration errors. Departures from normally-distributed errors can also increase  $\langle \chi^2 \rangle_\pi$ .

If  $D$  is not at fault, then  $\mathcal{M}$  is the culprit, implying that either  $\pi$  or  $H$  is at fault. If the fault lies with  $\pi$  not  $H$ , then we expect that  $\langle \chi^2 \rangle_u < \chi_{\nu, \beta}^2$  even though

$$\langle \chi^2 \rangle_\pi > \chi_{\nu, \beta}^2.$$

If neither  $D$  nor  $\pi$  can be faulted, then the investigator must seek a refined or alternative  $H$ .

#### 3.2. Type I and type II errors

In the frequentist approach to hypothesis testing, decision errors are said to be of type I if  $H$  is true but the test says reject  $H$ , and of type II if  $H$  is false but the test says accept  $H$ .

Since testing a Bayesian model is not concerned exclusively with  $H$ , these definitions must be revised, as follows:

- A type I error arises when  $\mathcal{M}$  and  $D$  are flawless but the statistic (e.g.,  $\langle \chi^2 \rangle_\pi$ ) exceeds the designated threshold.
- A type II error arise when  $\mathcal{M}$  or  $D$  are flawed but the statistic does not exceed the designated threshold.

Here the words accept and reject are avoided. Moreover, no particular threshold is mandatory: it is at the discretion of the investigator and is chosen with regard to the consequences of making a decision error.

## 4. Statistics of $\langle \chi^2 \rangle$

The intuitive understanding that scientists have regarding  $\chi_0^2$  derives from its simplicity and the rigorous theorems on its distribution that allow us to derive confidence regions for multi-dimensional linear models (e.g., Press et al. 2007, Sect. 15.6.4).

Rigorous statistics for  $\chi^2$  require two assumptions: 1) that the model fitted to  $D$  is linear in its parameters; and 2) that measurement errors obey the normal distribution. Nevertheless, even when these standard assumptions do not strictly hold, scientists still commonly rely on  $\chi_0^2$  to gauge goodness-of-fit, with perhaps Monte Carlo (MC) sampling to provide justification or calibration (e.g., Press et al. 2007, Sect. 15.6.1).

Rigorous results for the statistic  $\langle \chi^2 \rangle$  are therefore of interest. In fact, if we add the assumption of a uniform prior to the above standard assumptions, then we may prove (Appendix A) that

$$\langle \chi^2 \rangle_u = \chi_0^2 + k \quad (6)$$

where  $k$  is the number of parameters.

Given that Eq. (6) is exact under the stated assumptions, it follows that the quantity  $\langle \chi^2 \rangle_u - k$  is distributed as  $\chi_\nu^2$  with  $\nu = n - k$  degrees of freedom.

For minimum- $\chi^2$  fitting of linear models, the solution is always a better fit to  $D$  than is the true solution. In particular,  $\mathcal{E}(\chi_{\text{true}}^2) = n$ , whereas  $\mathcal{E}(\chi_0^2) = n - k$ . Accordingly, the  $+k$  term in Eq. (6) “corrects”  $\chi_0^2$  for overfitting and so  $\mathcal{E}(\langle \chi^2 \rangle_u) = \mathcal{E}(\chi_{\text{true}}^2)$  – i.e., the expected value of  $\chi^2$  for the actual measurement errors.

#### 4.1. Effect of an informative prior

When an informative prior is included in  $\mathcal{M}$ , an analytic formula for  $\langle \chi^2 \rangle_\pi$  is in general not available. However, its approximate statistical properties are readily found.

Consider again a linear model with normally-distributed errors and suppose further that the experiment ( $X_2$ ) is without flaws. The  $\chi^2$  surfaces are then self-similar  $k$ -dimensional ellipsoids with minimum at  $\alpha_0 \approx \alpha_{\text{true}}$ , the unknown exact solution. Let us now further suppose that the informative prior  $\pi(\alpha)$  derives from a previous flawless experiment ( $X_1$ ). The prior  $\pi$  will then be a convex (bell-shaped) function with maximum at  $\alpha_{\text{max}} \approx \alpha_0$ . Now, consider a 1D family of such functions all centred on  $\alpha_0$  and ranging from the very broad to the very narrow. For the former  $\langle \chi^2 \rangle_\pi \approx \langle \chi^2 \rangle_u$ ; for the latter  $\langle \chi^2 \rangle_\pi \approx \chi_0^2$ . Thus, ideally, when a Bayesian model  $\mathcal{M}$  is applied to data  $D$  we expect that

$$\langle \chi^2 \rangle_u \gtrsim \langle \chi^2 \rangle_\pi \geq \chi_0^2. \quad (7)$$

Now, a uniform prior and  $\delta(\alpha - \alpha_0)$  are the limits of the above family of bell-shaped functions. Since the delta function limit is not likely to be closely approached in practice, a first approximation to the distribution of  $\langle \chi^2 \rangle_\pi$  is that of  $\langle \chi^2 \rangle_u$  – i.e., that of  $\chi_{n-k}^2 + k$ .

The above discussion assumes faultless  $X_1$  and  $X_2$ . But now suppose that there is an inconsistency between  $\pi$  and  $X_2$ . The peak of  $\pi$  at  $\alpha_{\text{max}}$  will then in general be offset from the minimum of  $\chi^2$  at  $\alpha_0$ . Accordingly, in the calculation of  $\langle \chi^2 \rangle_\pi$  from Eq. (4), the neighbourhood of the  $\chi^2$  minimum  $\chi_0^2$  at  $\alpha_0$  has reduced weight relative to  $\chi^2(\alpha_{\text{max}})$  at the peak of  $\pi$ . Evidently, in this circumstance,  $\langle \chi^2 \rangle_\pi$  can greatly exceed  $\langle \chi^2 \rangle_u$ , and the investigator is then alerted to the inconsistency.

## 5. Numerical experiments

Given that rigorous results  $\langle \chi^2 \rangle$  are not available for informative priors, numerical tests are essential to illustrate the discussion of Sect. 4.1.

### 5.1. A toy model

A simple example with just one parameter  $\mu$  is as follows:  $H$  states that  $u = \mu$ , and  $D$  comprises  $n$  measurements  $u_i = \mu + \sigma z_i$ , where the  $z_i$  here and below are independent gaussian variates randomly sampling  $\mathcal{N}(0, 1)$ . In creating synthetic data, we set  $\mu = 0$ ,  $\sigma = 1$  and  $n = 100$ .

In the first numerical experiment, two independent data sets  $D_1$  and  $D_2$  are created comprising  $n_1$  and  $n_2$  measurements, respectively. On the assumption of a uniform prior, the posterior density of  $\mu$  derived from  $D_1$  is

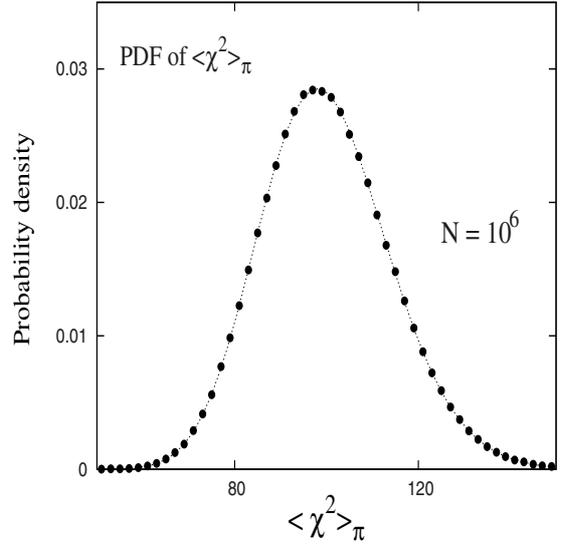
$$p(\mu|H, D_1) = \frac{\mathcal{L}(\mu|H, D_1)}{\int \mathcal{L}(\mu|H, D_1) d\mu}. \quad (8)$$

We now regard  $p(\mu|H, D_1)$  as prior knowledge to be taken into account in analysing  $D_2$ . Thus

$$\pi(\mu) = p(\mu|H, D_1) \quad (9)$$

so that the posterior distribution derived from  $D_2$  is

$$p(\mu|H, D_2) = \frac{\pi(\mu)\mathcal{L}(\mu|H, D_2)}{\int \pi(\mu)\mathcal{L}(\mu|H, D_2) d\mu} \quad (10)$$



**Fig. 1.** Empirical PDF of  $\langle \chi^2 \rangle_\pi$  derived from the analysis of  $10^6$  data pairs  $D_1, D_2$  as described in Sect. 5.1. The dashed curve is the theoretical PDF for  $\langle \chi^2 \rangle_u$ .

The statistic quantifying the goodness-of-fit achieved when  $\mathcal{M} = \{\pi, H\}$  is applied to data  $D_2$  is then

$$\langle \chi^2 \rangle_\pi = \int \chi^2(\mu) p(\mu|H, D_2) d\mu. \quad (11)$$

From  $N$  independent data pairs ( $D_1, D_2$ ), we obtain  $N$  independent values of  $\langle \chi^2 \rangle_\pi$  thus allowing us to test the expectation (Sect. 4.1) that this statistic is approximately distributed as  $\langle \chi^2 \rangle_u$ . In Fig. 1, this empirical PDF is plotted together the theoretical PDF for  $\langle \chi^2 \rangle_u$ . The accuracy of the approximation at large values of  $\chi^2$  is of special importance. For  $n = 100$  and  $k = 1$ , the 0.05, 0.01 and 0.001 critical points of  $\langle \chi^2 \rangle_u$  are 124.2, 135.6 and 149.2, respectively. From a simulation with  $N = 10^6$ , the number of  $\langle \chi^2 \rangle_\pi$  values exceeding these thresholds are 50 177, 10 011 and 1025, respectively. Thus, the fraction of  $\langle \chi^2 \rangle_\pi$  exceeding the critical values derived from the distribution of  $\langle \chi^2 \rangle_u$  are close to their predicted values. Accordingly, the conventional interpretation of these critical values is valid.

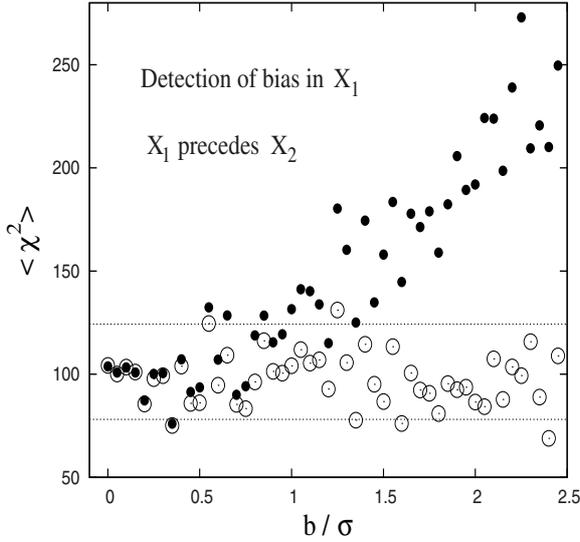
In this experiment, the analysis of  $X_2$  benefits from knowledge gained from  $X_1$ . We expect therefore that  $\langle \chi^2 \rangle_\pi$  is less than  $\langle \chi^2 \rangle_u$ , since replacing the uniform prior with the informative  $\pi$  obtained from  $X_1$  should improve the fit. From  $10^6$  repetitions, this proves to be so with probability 0.683. Sampling noise in  $D_1$  and  $D_2$  accounts for the shortfall.

### 5.2. Bias test

When  $X_1$  and  $X_2$  are flawless, the statistic  $\langle \chi^2 \rangle_\pi$  indicates doubts – i.e., type I errors (Sect. 3.2) – about the mutual consistency of  $X_1$  and  $X_2$  with just the expected frequency. Thus, with the 5% threshold, doubts arise in 5.02% of the above  $10^6$  trials. This encourages the use of  $\langle \chi^2 \rangle_\pi$  to detect inconsistency.

Accordingly, in a second test,  $X_1$  is flawed due to biased measurements. Thus, now  $u_i = \mu + \sigma z_i + b$  for  $D_1$ . As a result, the prior for  $X_2$  obtained from Eq. (9) is compromised, and this impacts on the statistic  $\langle \chi^2 \rangle_\pi$  from Eq. (11).

In Fig. 2, the values of  $\langle \chi^2 \rangle_\pi$  and  $\langle \chi^2 \rangle_u$  are plotted against  $b/\sigma$ . Since the compromised prior is excluded from  $\langle \chi^2 \rangle_u$ , its values depend only on the flawless data sets  $D_2$ , and



**Fig. 2.** Detection of bias in  $X_1$  with the  $\langle \chi^2 \rangle_\pi$  statistic when  $X_2$  is analysed with prior derived from  $X_1$ . Values of  $\langle \chi^2 \rangle_\pi$  (filled circles) and  $\langle \chi^2 \rangle_u$  (open circles) are plotted against the bias parameter  $b/\sigma$ . The dashed lines are the 5 and 95% levels.

so mostly fall within the (0.05, 0.95) interval. In contrast, as  $b/\sigma$  increases, the values of  $\langle \chi^2 \rangle_\pi$  are increasingly affected by the compromised prior.

The mutual consistency of  $X_1$  and  $X_2$  is assessed on the basis of  $\langle \chi^2 \rangle_\pi$ , with choice of critical level at our discretion. However, when  $\langle \chi^2 \rangle_\pi$  exceeds the 0.1% level at 149.2, we would surely conclude that  $X_1$  and  $X_2$  are in conflict and seek to resolve the discrepancy. On the other hand, when inconsistency is not indicated, we may accept the Bayesian inferences derived from  $X_2$  in the confident belief that incorporating prior knowledge from  $X_1$  is justified and beneficial. This test illustrates the important point that an inappropriate  $\pi$  can corrupt the Bayesian inferences drawn from a flawless experiment. Thus, in this case, the bias in  $D_1$  propagates into the posterior  $p(\mu|H, D_2)$  derived from  $X_2$ . This can be (and is) avoided by preferring the frequentist methodology. But to do so is to forgo the great merit of Bayesian inference, namely its ability to incorporate *informative* prior information (Feldman & Cousins 1998). If one does therefore prefer Bayesian inference, it is evident that a goodness-of-fit statistic such as  $\langle \chi^2 \rangle_\pi$  is essential in order to detect errors propagating into the posterior distribution from an ill-chosen prior.

### 5.3. Order reversed

In the above test, the analysis of  $X_2$  is preceded by that of  $X_1$ . This order can be reversed. Thus, with the same  $N$  data pairs  $(D_1, D_2)$ , we now first analyse  $X_2$  with a uniform prior to obtain  $p(\mu|H, D_2)$ . This becomes the prior for the analysis of  $X_1$ . This analysis then gives the posterior  $p(\mu|H, D_1)$  from which a new value of  $\langle \chi^2 \rangle_\pi$  is obtained.

When the values of  $\langle \chi^2 \rangle_\pi$  obtained with this reversed order of analysis are plotted against  $b/\sigma$ , the result is similar to Fig. 2, implying that the order is unimportant. Indeed, statistically, the same decision is reached independently of order. For example, for  $10^5$  independent data pairs  $(D_1, D_2)$  with  $b/\sigma = 1$ , the number with  $\langle \chi^2 \rangle_\pi > 124.2$ , the 5% threshold, is 50 267 when  $X_1$  precedes  $X_2$  and 50 149 when  $X_2$  precedes  $X_1$ .

### 5.4. Alternative statistic

Noting that the Bayesian evidence is  $= \bar{\mathcal{L}}$ , the prior-weighted mean of the likelihood, we can, under standard assumptions, write

$$\bar{\mathcal{L}} \propto \exp\left[-\frac{1}{2}\chi_{\text{eff}}^2\right] \quad (12)$$

where the effective  $\chi^2$  (Bridges et al. 2009) is

$$\chi_{\text{eff}}^2 = -2 \ln \int \pi(\alpha) \exp\left[-\frac{1}{2}\chi^2(\alpha)\right] dV_\alpha. \quad (13)$$

This is a possible alternative to  $\langle \chi^2 \rangle_\pi$  defined in Eq. (4). However, in the test of Sect. 5.1, the two values are so nearly identical it is immaterial which mean is used. Here  $\langle \chi^2 \rangle_\pi$  is preferred because it remains well-defined for a uniform prior, for which an analytic result is available (Appendix A).

Because  $\langle \chi^2 \rangle_\pi$  and  $\chi_{\text{eff}}^2$  are nearly identical, the distribution of  $\chi_{\text{eff}}^2$  should approximate that of  $\langle \chi^2 \rangle_u$  (Sect. 4.1). To test this, the experiment of Sect. 5.1 is repeated with  $\chi_{\text{eff}}^2$  replacing  $\langle \chi^2 \rangle_\pi$ . From a simulation with  $N = 10^6$ , the number of  $\chi_{\text{eff}}^2$  values exceeding the 0.05, 0.01 and 0.001 thresholds are 50 167, 9951 and 970, respectively. Thus if  $\chi_{\text{eff}}^2$  is chosen as the goodness-of-fit statistic, accurate  $p$ -values can be derived on the assumption that  $\chi_{\text{eff}}^2 - k$  is distributed as  $\chi_\nu^2$  with  $\nu = n - k$  degrees of freedom. From Sect. 4.1, we expect these  $p$ -values to be accurate if  $\pi(\alpha)$  is *not* more sharply peaked than  $\mathcal{L}(\alpha|H, D)$ .

## 6. An $F$ statistic for Bayesian models

Inspection of Fig. 2 shows that a more powerful test of inconsistency between  $X_1$  and  $X_2$  must exist. A systematic displacement of  $\langle \chi^2 \rangle_\pi$  relative to  $\langle \chi^2 \rangle_u$  is already evident even when  $\langle \chi^2 \rangle_\pi$  is below the 5% threshold at 124.2. This suggests that a Bayesian analogue of the  $F$  statistic be constructed.

### 6.1. The frequentist $F$ -test

In frequentist statistics, a standard result (e.g., Hamilton 1964, p. 139) in the testing of *linear* hypotheses is the following: we define the statistic

$$F = \frac{n-i}{j} \frac{\chi_c^2 - \chi_0^2}{\chi_0^2} \quad (14)$$

where  $\chi_0^2$  is the minimum value of  $\chi^2$  when all  $i$  parameters are adjusted, and  $\chi_c^2$  is the minimum value when a linear constraint is imposed on  $j$  ( $\leq i$ ) parameters, so that only the remaining  $i - j$  are adjusted. Then, on the null hypothesis  $H$  that the constraint is true,  $F$  is distributed as  $F_{\nu_1, \nu_2}$  with  $\nu_1 = j$  and  $\nu_2 = n - i$ , where  $n$  is the number of measurements. Accordingly,  $H$  is tested by comparing the value  $F$  given by Eq. (14) with critical values  $F_{\nu_1, \nu_2, \beta}$  derived from the distribution of  $F_{\nu_1, \nu_2}$ .

Note that when  $j = i$ , the constraint completely determines  $\alpha$ . If this value is  $\alpha_*$ , then  $\chi_c^2 = \chi^2(\alpha_*)$  and  $H$  states that  $\alpha_* = \alpha_{\text{true}}$ .

A particular merit of the statistic  $F$  is that it is independent of  $\sigma$ . However, the resulting  $F$ -test does assume normally-distributed measurement errors.

## 6.2. A Bayesian $F$

In a Bayesian context, the frequentist hypothesis that  $\alpha_{\text{true}} = \alpha_*$  is replaced by the statement that  $\alpha_{\text{true}}$  obeys the posterior distribution  $p(\alpha|H, D_2)$ . Thus an exact constraint is replaced by a fuzzy constraint.

Adopting the simplest approach, we define  $\mathcal{F}$ , a Bayesian analogue of  $F$ , by taking  $\chi_c^2$  to be the value at the posterior mean of  $\alpha$ ,

$$\langle \alpha \rangle_\pi = \frac{\int \alpha \pi(\alpha) \mathcal{L}(\alpha|H, D_2) dV_\alpha}{\int \pi(\alpha) \mathcal{L}(\alpha|H, D_2) dV_\alpha} \quad (15)$$

where  $\pi(\alpha)$  is the informative prior. Considerations of accuracy when values of  $\chi^2$  are computed on a grid suggest we take  $\chi_0^2$  to be the value at

$$\langle \alpha \rangle_u = \frac{\int \alpha \mathcal{L}(\alpha|H, D_2) dV_\alpha}{\int \mathcal{L}(\alpha|H, D_2) dV_\alpha} \quad (16)$$

the posterior mean for a uniform prior.

With  $\chi_c^2$  and  $\chi_0^2$  thus defined, the Bayesian  $F$  is

$$\mathcal{F} = \frac{n-i}{j} \frac{\chi^2(\langle \alpha \rangle_\pi) - \chi^2(\langle \alpha \rangle_u)}{\chi^2(\langle \alpha \rangle_u)} \quad (17)$$

and its value is to be compared with the chosen threshold  $F_{\nu_1, \nu_2, \beta}$  when testing the consistency of  $X_1$  and  $X_2$ . Since  $\langle \alpha \rangle_u$  is independent of  $\pi(\alpha)$ , the statistic  $\mathcal{F}$  measures the effect of  $\pi(\alpha)$  in displacing  $\langle \alpha \rangle_\pi$  from  $\langle \alpha \rangle_u$ .

In this Bayesian version of the  $F$ -test, the null hypothesis states that the posterior mean  $\langle \alpha \rangle_\pi = \alpha_{\text{true}}$ . This will be approximately true when a flawless Bayesian model  $\mathcal{M}$  is applied to a flawless data set  $D$ . However, if the chosen threshold on  $\mathcal{F}$  is exceeded, then one or more of  $\pi, H$  and  $D$  is suspect as discussed in Sect. 3.1. If the threshold is not exceeded, then Bayesian inferences drawn from the posterior distribution  $p(\alpha|H, D_2)$  are supported.

## 6.3. Test of $p$ -values from $\mathcal{F}$

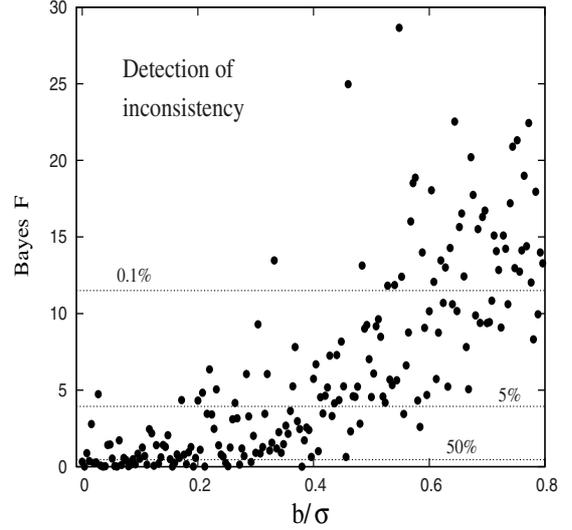
If Eq. (17) gives  $\mathcal{F} = \mathcal{F}_*$ , the corresponding  $p$ -value is

$$p_* = \int_{\mathcal{F}_*}^{\infty} F_{\nu_1, \nu_2} d\mathcal{F} \quad (18)$$

where  $\nu_1 = j$  and  $\nu_2 = n - i$ . The accuracy of these  $p$ -values can be tested with the 1D toy model of Sect. 5.1 as follows:

- (i) Independent data sets  $D_1, D_2$  are created corresponding to  $X_1, X_2$ .
- (ii) With  $\pi$  from  $X_1$ , the quantities  $\langle \mu \rangle_\pi^*$  and  $\mathcal{F}_*$  are calculated for  $X_2$  with Eqs. (15)–(17).
- (iii)  $M$  independent data sets  $\mathcal{D}_m$  are now created with  $u_i = \langle \mu \rangle_\pi^* + \sigma z_i$ .
- (iv) For each  $\mathcal{D}_m$ , the new value of  $\chi^2(\langle \alpha \rangle_\pi^*)$  is calculated for the  $\langle \mu \rangle_\pi^*$  derived at step (ii).
- (v) For each  $\mathcal{D}_m$ , the new value of  $\chi^2(\langle \alpha \rangle_u)$  is calculated with  $\langle \alpha \rangle_u$  from Eq. (16).
- (vi) With these new  $\chi^2$  values, the statistic  $\mathcal{F}_m$  is obtained from Eq. (17).

The resulting  $M$  values of  $\mathcal{F}_m$  give us an empirical estimate of  $p_*$ , namely  $f_*$ , the fraction of the  $\mathcal{F}_m$  that exceed  $\mathcal{F}_*$ . In one example of this test, a data pair  $D_1, D_2$  gives  $\langle \mu \rangle_\pi^* = 0.089$  and  $\mathcal{F}_* = 1.0021$ . With  $\nu_1 = 1$  and  $\nu_2 = 99$ , Eq. (18) then gives



**Fig. 3.** Detection of inconsistency between  $X_1$  and  $X_2$ . Values of  $\mathcal{F}$  are plotted against the bias parameter  $b/\sigma$ . The dashed lines are the 0.1, 5 and 50% levels. Note that the horizontal scale differs from Fig. 2.

$p_* = 0.3192$ . This is checked by creating  $M = 10^5$  data sets  $\mathcal{D}_m$  with the assumption that  $\mu_{\text{true}} = 0.089$ . The resulting empirical estimate is  $f_* = 0.3189$ , in close agreement with  $p_*$ .

From 100 independent pairs  $D_1, D_2$ , the mean value of  $|p_* - f_*|$  is 0.001. This simulation confirms the accuracy of  $p$ -values derived from Eq. (18) and therefore of decision thresholds  $F_{\nu_1, \nu_2, \beta}$ .

## 6.4. Bias test

To investigate this Bayesian  $F$ -test's ability to detect inconsistency between  $X_1$  and  $X_2$ , the bias test of Sect. 5.2 is repeated, again with  $n_1 = n_2 = 100$ . The vector  $\alpha$  in Sect. 6.1 now becomes the scalar  $\mu$ , and  $i = 1$ .

In Fig. 3, the values of  $\mathcal{F}$  from Eq. (17) with  $j = i = 1$  are plotted against the bias parameter  $b/\sigma$ . Also plotted are critical values derived from the distribution  $F_{\nu_1, \nu_2}$  with  $\nu_1 = 1$  and  $\nu_2 = 99$ . In contrast to Fig. 2 for the  $\langle \chi^2 \rangle_\pi$  statistic, inconsistency between  $X_1$  and  $X_2$  is now detected down to  $b/\sigma \approx 0.5$ .

In this test of inconsistency, the flaw in  $X_1$  is the bias  $b$ . Now, if it were known for certain that this was the flaw, then a Bayesian analysis with  $H_1$  changed from  $u = \mu$  to  $u = \mu + b$  – i.e., with an extra parameter – is straightforward. The result is the posterior density for  $b$ , allowing for correction. In contrast, the detection of a flaw with  $\langle \chi^2 \rangle_\pi$  or  $\mathcal{F}$  is cause-independent. Although Figs. 2 and 3 have  $b/\sigma$  as the abscissa, for real experiments this quantity is not known and conclusions are drawn just from the ordinate.

## 7. Tension between experiments

In the above, the goodness-of-fit of  $\mathcal{M}$  to  $D$  is investigated taking into account the possibility that a poor fit might be due the prior derived from a previous experiment. A related goodness-of-fit issue commonly arises in modern astrophysics, particularly cosmology, namely whether estimates of non-identical but partially overlapping parameter vectors from different experiments are mutually consistent. The term commonly used in this regard is tension, with investigators often reporting their *subjective* assessments – e.g., there is marginal tension between  $X_1$  and  $X_2$  – based on the two credibility domains (often multi-dimensional) for the parameters in common.

In a recent paper, Seehars et al. (2015) review the attempts in the cosmological literature to quantify the concept of tension, with emphasis on CMB experiments. Below, we develop a rather different approach based on the  $\mathcal{F}$  statistic defined in Sect. 6.2.

Since detecting and resolving conflicts between experiments is essential for scientific progress, it is desirable to quantify the term tension and to optimize its detection. The conjecture here is that this optimum is achieved when inferences from  $X_1$  are imposed on the analysis of  $X_2$ .

### 7.1. Identical parameter vectors

A special case of assessing tension between experiments is that where the parameter vectors are identical. But this is just the problem investigated in Sects. 5 and 6.

When  $X_1$  (with bias) and  $X_2$  (without bias) are separately analysed, the limited overlap of the two credibility intervals for  $\mu$  provides a *qualitative* indication of tension. However, if  $X_2$  is analysed with a prior derived from  $X_1$ , then the statistic  $\langle \chi^2 \rangle_\pi$  – see Fig. 2 – or, more powerfully, the statistic  $\mathcal{F}$  – see Fig. 3 – provide a *quantitative* measure to inform statements about the degree of tension.

### 7.2. Non-identical parameter vectors-I

We now suppose that  $D_1, D_2$  are data sets from different experiments  $X_1, X_2$  designed to test the hypotheses  $H_1, H_2$ . However, though different, these hypotheses have parameters in common. Specifically, the parameter vectors of  $H_1$  and  $H_2$  are  $\xi = (\alpha, \beta)$  and  $\eta = (\beta, \gamma)$ , respectively, and  $k, l, m$  are the numbers of elements in  $\alpha, \beta, \gamma$ , respectively.

If  $X_1$  and  $X_2$  are analysed independently, we may well find that  $\mathcal{M}_1$  and  $\mathcal{M}_2$  provide satisfactory fits to  $D_1$  and  $D_2$  and yet still be obliged to report tension between the experiments because of a perceived inadequate overlap of the two  $l$ -dimensional credibility domains for  $\beta$ .

A quantitative measure of tension between  $X_1$  and  $X_2$  can be derived via the priors, as follows: the analysis of  $X_1$  gives  $p(\xi|H_1, D_1)$ , the posterior distribution of  $\xi$ , from which we may derive the posterior distribution of  $\beta$  by integrating over  $\alpha$ . Thus

$$p(\beta|H_1, D_1) = \int p(\xi|H_1, D_1) dV_\alpha. \quad (19)$$

Now, for a Bayesian analysis of  $X_2$ , we must specify  $\pi(\eta)$  throughout  $\eta$ -space, not just  $\beta$ -space. But

$$\pi(\eta) = \pi(\beta, \gamma) = \pi(\gamma|\beta) \pi(\beta). \quad (20)$$

Accordingly, what we infer from  $X_1$  can be imposed on the analysis of  $X_2$  by writing

$$\pi(\eta) = \pi(\gamma|\beta) p(\beta|H_1, D_1). \quad (21)$$

The conditional prior  $\pi(\gamma|\beta)$  must now be specified. This can be taken to be uniform unless we have prior knowledge from other sources – i.e., not from  $D_2$ .

With  $\pi(\eta)$  specified, the analysis of  $X_2$  gives the posterior density  $p(\eta|H_2, D_2)$ . As in Sect. 6.2, we now regard this as a fuzzy constraint on  $\eta$  from which we compute the sharp constraint  $\langle \eta \rangle_\pi$  given by Eq. (15). Now, in general,  $\langle \eta \rangle_\pi$  will be displaced from  $\langle \eta \rangle_u$  given by Eq. (16). The question then is: Is the increment in  $\chi^2(\eta|H_2, D_2)$  between  $\langle \eta \rangle_\pi$  and  $\langle \eta \rangle_u$  so large that we must acknowledge tension between  $X_1$  and  $X_2$ ?

Following Sect. 6, we answer this question by computing  $\mathcal{F}$  from Eq. (17) with  $i = j = l + m$ , the total number parameters in  $\eta$ . The result is then compared to selected critical values from the  $F_{\nu_1, \nu_2}$  distribution, where  $\nu_1 = l + m$  and  $\nu_2 = n_2 - l - m$ . With standard assumptions,  $\mathcal{F}$  obeys this distribution if  $\langle \eta \rangle_\pi = \eta_{\text{true}}$  – i.e., if  $\langle \beta \rangle_\pi = \beta_{\text{true}}$  and  $\langle \gamma \rangle_\pi = \gamma_{\text{true}}$  – see Sect. 6.3.

### 7.3. A toy model

A simple example with one parameter ( $\mu$ ) for  $X_1$  and two ( $\mu, \kappa$ ) for  $X_2$  is as follows:  $H_1$  states that  $u = \mu$  and  $H_2$  states that  $v = \mu + \kappa x$ . The data set  $D_1$  comprises  $n_1$  measurements  $u_i = \mu + \sigma z_i + b$ , where  $b$  is the bias. The data set  $D_2$  comprises  $n_2$  measurements  $v_j = \mu + \kappa x_j + \sigma z_j$ , where the  $x_j$  are uniform in  $(-1, +1)$ . The parameters are  $\mu = 0, \kappa = 1, \sigma = 1$  and  $n_1 = n_2 = 100$ . In the notation of Sect. 7.2, the vectors  $\beta, \gamma$  contract to the scalars  $\mu, \kappa$ , whence  $l = m = 1$ , and  $\alpha$  does not appear, whence  $k = 0$ .

### 7.4. Bias test

In the above,  $H_1$  and  $H_2$  are different hypotheses but have the parameter  $\mu$  in common. If  $b = 0$ , the analyses of  $X_1$  and  $X_2$  should give similar credibility intervals for  $\mu$  and therefore no tension. But with sufficiently large  $b$ , tension should arise.

This is investigated following Sect. 7.2. Applying  $\mathcal{M}_1$  to  $D_1$ , we derive  $p(\mu|H_1, D_1)$ . Then, taking the conditional prior  $\pi(\kappa|\mu)$  to be constant, we obtain

$$\pi(\mu, \kappa) \propto p(\mu|H_1, D_1) \quad (22)$$

as the prior for the analysis of  $X_2$ . This gives us the posterior distribution  $p(\mu, \kappa|H_2, D_2)$ , which is a fuzzy constraint in  $(\mu, \kappa)$ -space. Replacing this by the sharp constraint  $(\langle \mu \rangle, \langle \kappa \rangle)$ , the constrained  $\chi^2$  is

$$\chi_c^2 = \chi^2(\langle \mu \rangle, \langle \kappa \rangle|H_2, D_2). \quad (23)$$

Substitution in Eq. (17) with  $j = i = 2$ , then gives  $\mathcal{F}$  as a measure of the tension between  $X_1$  and  $X_2$ . Under standard assumptions,  $\mathcal{F}$  is distributed as  $F_{\nu_1, \nu_2}$  with  $\nu_1 = 2, \nu_2 = n_2 - 2$  if  $(\langle \mu \rangle, \langle \kappa \rangle) = (\mu, \kappa)_{\text{true}}$ .

In Fig. 4, the values of  $\mathcal{F}$  are plotted against  $b/\sigma$  together with critical values for  $F_{\nu_1, \nu_2}$  with  $\nu_1 = 2, \nu_2 = 98$ . This plot shows that tension is detected for  $b/\sigma \gtrsim 0.6$ . This is slightly inferior to Fig. 3 as is to be expected because of the more complicated  $X_2$ .

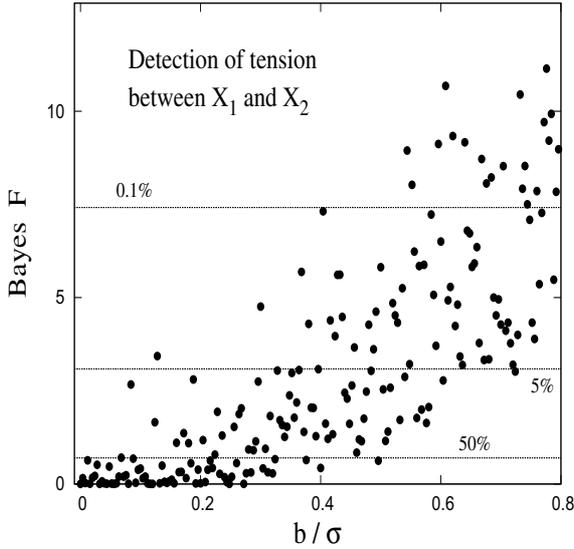
The importance of Fig. 4 is in showing that the  $\mathcal{F}$  statistic has the desired characteristic of reliably informing the investigator of tension between different experiments with partially overlapping parameter vectors. When the inconsistency is slight ( $b/\sigma \lesssim 0.2$ ), this statistic does not sound a false alarm. When the inconsistency is substantial ( $b/\sigma \gtrsim 0.6$ ), the statistic does not fail to sound the alarm.

### 7.5. Non-identical parameter vectors-II

If  $\mathcal{F}$  calculated as in Sect. 7.2 indicates tension, the possible flaws include the conditional prior  $\pi(\gamma|\beta)$ . Thus, tension could be indicated even if the prior  $\pi(\beta)$  inferred from  $X_1$  is accurate and consistent with  $D_2$ .

Accordingly, we might prefer to focus just on  $\beta$  – i.e., on the parameters in common. If so, we compute

$$\langle \beta \rangle_\pi = \int \beta p(\eta|H_2, D_2) dV_\eta \quad (24)$$



**Fig. 4.** Detection of tension between different experiments. Values of  $\mathcal{F}$  are plotted against the bias parameter  $b/\sigma$ . The dashed lines are the 0.1, 5 and 50% levels.

and then find the minimum of  $\chi^2(\eta|H_2, D_2)$  when  $\beta = \langle\beta\rangle_\pi$ . Thus, the constrained  $\chi^2$  is now

$$\chi_c^2 = \min_{\gamma} \chi^2(\langle\beta\rangle_\pi, \gamma). \quad (25)$$

The  $F$  test also applies in this case – see Sect. 6.1. Thus this value  $\chi_c^2$  is substituted in Eq. (17), where we now take  $j = l$ , the number of parameters in  $\beta$ . The resulting  $\mathcal{F}$  is then compared to critical values derived from  $F_{\nu_1, \nu_2}$  with  $\nu_1 = l$ ,  $\nu_2 = n_2 - l - m$ . With the standard assumptions,  $\mathcal{F}$  obeys this distribution if  $\langle\beta\rangle_\pi = \beta_{\text{true}}$ .

For the simple model of Sect. 7.3, the resulting plot of  $\mathcal{F}$  against  $b/\sigma$  is closely similar to Fig. 4 and so is omitted. Nevertheless, the option of constraining a subset of the parameters is likely to be a powerful diagnostic tool for complex, multi-dimensional problems, identifying the parameters contributing most to the tension revealed when the entire vector is constrained (cf. Seehars et al. 2015).

## 8. Discussion and conclusion

A legitimate question to ask about the statistical analysis of data acquired in a scientific experiment is: How well or how badly does the model fit the data? Asking this question is, after all, just the last step in applying the scientific method. In a frequentist analysis, where the estimated parameter vector  $\alpha_0$  is typically the minimum- $\chi^2$  point, this question is answered by reporting the goodness-of-fit statistic  $\chi_0^2 = \chi^2(\alpha_0)$  or, equivalently, the corresponding  $p$ -value. If a poor fit is thereby indicated, the investigator and the community are aware that not only is the model called into question but so also is the estimate  $\alpha_0$  and its confidence domain.

If the same data is subject to a Bayesian analysis, the same question must surely be asked: the Bayesian approach does not exempt the investigator from the obligation to report on the success or otherwise of the adopted model. In this case, if the fit is poor, the Bayesian model is called into question and so also are inferences drawn from the posterior distribution  $p(\alpha|H, D)$ .

As noted in Sect. 1, the difficulty in testing Bayesian models is that there are no good fully Bayesian methods for assessing goodness-of-fit. Accordingly, in this paper, a frequentist approach is advocated. Specifically,  $\langle\chi^2\rangle_\pi$  is proposed in Sect. 3 as a suitable goodness-of-fit statistic for Bayesian models. Under the null hypothesis that  $\mathcal{M}$  and  $D$  are flawless and with the standard assumptions of linearity and normally-distributed errors, then, as argued in Sect. 4.1 and illustrated in Fig. 1,  $\langle\chi^2\rangle_\pi - k$  is approximately distributed as  $\chi_{n-k}^2$ , and so  $p$ -values can be computed. A  $p$ -value thus derived from  $\langle\chi^2\rangle_\pi$  quantifies the average goodness-of-fit provided by the posterior distribution. In contrast, in a frequentist minimum- $\chi^2$  analysis, the  $p$ -value quantifies the goodness-of-fit provided by the point estimate  $\alpha_0$ .

In the above, it is regarded as self-evident that astronomers want to adhere to the scientific method by always reporting the goodness-of-fit achieved in Bayesian analyses of observational data. However, Gelman & Shalizi (2013), in an essay on the philosophy and practice of Bayesian statistics, note that investigators who identify Bayesian inference with inductive inference regularly fit and compare models without checking them. They deplore this practice. Instead, these authors advocate the hypothetico-deductive approach in which model checking is crucial. As in this paper, they discuss *non-Bayesian* checking of Bayesian models – specifically, the derivation of  $p$ -values from posterior predictive distributions. Moreover, they also stress that the prior distribution is a testable part of a Bayesian model.

In the astronomical literature, the use of frequentist tests to validate Bayesian models is not unique to this paper. Recently, Harrison et al. (2015) have presented an ingenious procedure for validating multidimensional posterior distributions with the frequentist Kolmogorov-Smirnov (KS) test for one-dimensional data. Their aim, as here, is to test the entire Bayesian inference procedure.

Frequentist testing also arises in recent applications of the Kullback-Leibler divergence to quantify tension between cosmological probes (e.g. Seehars et al. 2015). For linear models, and with the assumption of Gaussian priors and likelihoods, a term in the relative entropy is a statistic that measures tension. With these assumptions, the statistic follows a generalized  $\chi^2$  distribution, thus allowing a  $p$ -value to be computed.

Seehars et al. (2015) also investigate various purely Bayesian measures of tension. They conclude that interpreting these measures on a fixed, problem-independent scale – e.g., the Jeffrey’s scale – can be misleading – see also Nesseris & García-Bellido (2013).

*Acknowledgements.* I thank D. J. Mortlock for comments on an early draft of this paper, and A. H. Jaffe, M.P.Hobson and the referee for useful references.

## References

- Bridges, M., Feroz, F., Hobson, M. P., & Lasenby, A. N. 2009, *MNRAS*, **400**, 1075
- Feldman, G. J., & Cousins, R. D. 1998, *Phys. Rev. D*, **57**, 3873
- Gelman, A., & Shalizi 2013, *British J. Math. Stat. Psychol.*, **66**, 8
- Hamilton, W. C. 1964, *Statistics in Physical Science* (New York: Ronald Press)
- Harrison, D., Sutton, D., Carvalho, P., & Hobson, M. 2015, *MNRAS*, **451**, 2610
- Lucy, L. B. 2014, *A&A*, **571**, A86 (L14)
- Nesseris, S., & García-Bellido, J. 2013, *J. Cosmol. Astropart. Phys.*, **08**, 036
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007 *Numerical Recipes* (Cambridge: Cambridge University Press)
- Seehars, S., Grandis, S., Amara, A., & Refregier, A. 2015, arXiv e-prints [[arXiv:1510.08483](https://arxiv.org/abs/1510.08483)]
- Sellke, T., Bayarri, M. J., & Berger, J. 2001, *The American Statistician*, **55**, 62
- Sivia, D. S., & Skilling, J. 2006, *Data Analysis: A Bayesian Tutorial*, 2nd Edn. (Oxford: Oxford University Press)

## Appendix A: Evaluation of $\langle \chi^2 \rangle_u$

If  $\alpha_0$  denotes the minimum- $\chi^2$  solution, then

$$\alpha = \alpha_0 + \mathbf{a} \quad (\text{A.1})$$

where  $\mathbf{a}$  is the displacement from  $\alpha_0$ . Then, on the assumption of linearity,

$$\Delta\chi^2(\mathbf{a}) = \chi_0^2 + \sum_{i,j} A_{ij} a_i a_j \quad (\text{A.2})$$

where the  $A_{ij}$  are the constant elements of the  $k \times k$  curvature matrix (Press et al. 2007, p. 680), where  $k$  is the number of parameters. It follows that surfaces of constant  $\chi^2$  are  $k$ -dimensional self-similar ellipsoids centered on  $\alpha_0$ .

Now, given the second assumption of normally-distributed measurement errors, the likelihood

$$\mathcal{L}(\alpha) \propto \exp\left(-\frac{1}{2}\chi_0^2\right) \times \exp\left(-\frac{1}{2}\Delta\chi^2\right). \quad (\text{A.3})$$

Thus, in the case of a uniform prior, the posterior mean of  $\chi^2$  is

$$\langle \chi^2 \rangle_u = \chi_0^2 + \frac{\int \Delta\chi^2 \exp(-\frac{1}{2}\Delta\chi^2) dV_\alpha}{\int \exp(-\frac{1}{2}\Delta\chi^2) dV_\alpha}. \quad (\text{A.4})$$

Because surfaces of constant  $\Delta\chi^2$  are self-similar, the  $k$ -dimensional integrals in Eq. (A.4) reduce to 1D integrals.

Suppose  $\Delta\chi^2 = \Delta\chi_*^2$  on the surface of the ellipsoid with volume  $V_*$ . If lengths are increased by the factor  $\lambda$ , then the new ellipsoid has

$$\Delta\chi^2 = \Delta\chi_*^2 \times \lambda^2 \quad \text{and} \quad V = V_* \times \lambda^k. \quad (\text{A.5})$$

With these scaling relations, the integrals in Eq. (A.4) can be transformed into integrals over  $\lambda$ . The result is

$$\langle \chi^2 \rangle_u = \chi_0^2 + 2b \frac{\int_0^\infty \lambda^{k+1} \exp(-b\lambda^2) d\lambda}{\int_0^\infty \lambda^{k-1} \exp(-b\lambda^2) d\lambda} \quad (\text{A.6})$$

where  $2b = \Delta\chi_*^2$ . The integrals have now been transformed to a known definite integral,

$$\int_0^\infty \lambda^z \exp(-b\lambda^2) d\lambda = \frac{1}{2} \Gamma(x) b^{-x} \quad (\text{A.7})$$

where  $x = (z + 1)/2$ . Applying this formula, we obtain

$$\langle \chi^2 \rangle_u = \chi_0^2 + k \quad (\text{A.8})$$

an exact result under the stated assumptions.