Astronomy
&
Astrophysics

Binary orbits from combined astrometric and spectroscopic data

L. B. Lucy[*]

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

ABSTRACT

An efficient Bayesian technique for estimation problems in fundamental stellar astronomy is tested on simulated data for a binary observed both astrometrically and spectroscopically. Posterior distributions are computed for the components' masses and for the binary's parallax. One thousand independent repetitions of the simulation demonstrate that the 1- and 2-σ credibility intervals for these fundamental quantities have close to the correct coverage fractions. In addition, the simulations allow the investigation of the statistical properties of a Bayesian goodness-of-fit criterion and of the corresponding p-value. The criterion has closely similar properties to the traditional χ^2 test for minimum-χ^2 solutions.

Key words. binaries: visual – binaries: spectroscopic – stars: fundamental parameters – methods: statistical

1. Introduction

In fundamental stellar astronomy, all statistical estimation problems involve mathematical models with both linear and non-linear parameters – the so-called hybrid problems. The linear parameters determine scale and location; the non-linear parameters appear as arguments in dimensionless functions of time.

The presence of linearities suggests that more efficient estimation techniques exist than when all parameters are treated as non-linear (Wright & Howard 2009; Catanzarite 2010). However, some attempts to achieve this lead to significantly underestimated error bars (Eastman et al. 2013; Lucy 2014a, hereafter L14a). This poses the challenge of developing a technique that achieves the computational efficiency allowed by linearity while still giving confidence or credibility intervals with correct coverage – so that, for example, 1-σ error bars contain the correct answer with probability 0.683. A solution to this challenge is presented in (Lucy 2014b, hereafter L14b) where a grid search in the space defined by the non-linear parameters is combined with Monte Carlo sampling of the space defined by the linear parameters. In L14b, this technique is applied to simulated observations of a visual binary and coverage experiments confirm that 1- and 2-σ error bars enclose the exact values with close to the frequencies expected for normally distributed errors.

In this paper, a significantly harder problem is posed, that of analysing a binary with both astrometric and spectroscopic data. Such data could be analysed separately, but this is sub-optimal since information concerning several orbital elements is present in both data sets. Accordingly, the aim here is to obtain the posterior distribution over the entire parameter space using both data sets and to test if the derived error bars are trustworthy, an essential requirement for fundamental data on stellar masses and luminosities.

The posed problem aims at demonstrating proof-of-concept in the treatment of hybrid problems and to provide a template for

the many such problems in statistical astronomy. To this end, the code developed for this investigation is freely available.

Although here a technical exercise, the simultaneous analysis of astrometric and spectroscopic data is of practical importance in the era of adaptive optics (AO) and speckle interferometry. As emphasized by Mason et al. (1999), the ability to resolve binary stars at or near the diffraction limit results in a powerful synergy between short-period visual and long-period spectroscopic binaries, leading to stellar masses and improved mass-luminosity relations.

2. Synthetic orbits

The physical model comprises an isolated pair of stars undergoing Keplerian motion due to their mutual gravitational attraction. This binary is observed astrometrically and spectroscopically, yielding two independent data sets D_a and D_s, respectively. To analyse these data sets, the mathematical models predicting the components' relative motion on the sky as well as their radial velocity variations are used simultaneously to derive the posterior distribution over parameter space.

2.1. Orbital elements

For the astrometric orbit, L14a is followed closely with regard both to notation and the creation of synthetic data.

The motion on the sky of the secondary relative to the primary is conventionally parameterized by the Campbell elements $P, e, T, a, i, \omega, \Omega$. Here P is the period, e is the eccentricity, T is a time of periastron passage, i is the inclination, ω is the longitude of periastron, and Ω is the position angle of the ascending node. However, from the standpoint of computational economy, many investigators – references in L14a, Sect. 2.1 – prefer the Thiele-Innes elements. Thus, the Campbell parameter vector $\theta = (\phi, \vartheta)$, where $\phi = (P, e, \tau)$ and $\vartheta = (a, i, \omega, \Omega)$, is replaced by the Thiele-Innes vector (ϕ, ψ), where the components of the vector ψ are the Thiele-Innes constants A, B, F, G (Note that in the ϕ vector, T has been replaced by $\tau = T/P$ which by definition $\in (0, 1)$).

The spectroscopic orbits of the components introduce three additional parameters, the systemic velocity γ and the semi-amplitudes $K_{1,2}$. The predicted radial velocities are then

$$v_{1,2}(t) = \gamma + K_{1,2}\,[cos(\nu + \omega_{1,2}) + e\,cos\,\omega_{1,2}], \tag{1}$$

where $\nu(t)$ is the true anomaly, $\omega_2 = \omega$ and $\omega_1 = \omega + \pi$.

Note that $v_2 - v_1 = \dot{z}$, where z, the companion's displacement perpendicular to the sky, is given by the Thiele-Innes constants C and H. This is a useful check on coding.

With the inclusion of spectroscopic data, the combined data sets allow inferences about the 10-dimensional vector

$$\Theta = (\phi, \psi, \lambda), \tag{2}$$

where $\lambda = (\gamma, K_1, K_2)$.

In a Bayesian analysis, the task is to compute the posterior probability density in Θ-space given D_a and D_s.

2.2. Model binary

The adopted model binary has the following Campbell elements:

$$\begin{aligned} P = 10y \quad \tau = 0.4 \quad e = 0.6 \quad a = 0.3'' \\ i = 70° \quad \omega = 250° \quad \Omega = 120°, \end{aligned} \tag{3}$$

With this a, the binary would be unresolved in seeing-broadened images but should be resolved in images approaching diffraction limits.

If we now take the parallax $\varpi = 0.05''$, the total mass $= 2.16\,\mathcal{M}_\odot$. With mass ratio $q = 0.7$, the component masses are $\mathcal{M}_1 = 1.27$ and $\mathcal{M}_2 = 0.889\,\mathcal{M}_\odot$. The resulting semi-amplitudes are $K_1 = 8.64$ and $K_2 = 12.35\,km\,s^{-1}$, and without loss of generality we take $\gamma = 0.0\,km\,s^{-1}$.

2.3. Observing campaigns

An astrometric observing campaign is simulated by creating measured Cartesian sky coordinates $(\widetilde{x}_n, \widetilde{y}_n)$ with weights $w_n^a = 1/\sigma_{a,n}^2$ for both coordinates. One observation in observing seasons of length 0.3 yr is created randomly for 10 successive years. We take $\sigma_{a,n} = 0.05''$ for all n.

In the above, we assume equal precision for each coordinate *and* uncorrelated errors. These assumptions are well justified if the AO or speckle image reconstructions give circularly symmetric stellar profiles. If necessary, the technique can be generalized to treat unequal and correlated errors (Appendices A.1, C.1).

A spectroscopic observing campaign is simulated by creating measured radial velocities $\widetilde{v}_{1n}, \widetilde{v}_{2n}$ at random times in 10 successive observing seasons. The observations have weights $w_{1n}^s = 1/\sigma_{s,n}^2$ and $w_{2n}^s = 0.5/\sigma_{s,n}^2$. We take $\sigma_{s,n} = 0.5\,km\,s^{-1}$ for all n.

All 40 simulated observations have errors that are normally-distributed.

3. Conditional probabilities

In order to benefit from the hybrid character of the problems arising in orbit estimation, the chain rule for conditional probabilities is used to factorize multi-dimensional posterior distributions Λ in such a way that linear and non-linear parameters are separated. This facilitates the construction of efficient hybrid numerical schemes that combine grid scanning with Monte Carlo sampling.

3.1. Approach

Consider a problem with two scalar parameters, α and β, and suppose the model is non-linear in α and linear in β. Applying the chain rule, we can write the posterior density as

$$\Lambda(\alpha, \beta) = Pr(\alpha)\,Pr(\beta|\alpha), \tag{4}$$

where

$$Pr(\alpha) = \int \Lambda(\alpha, \beta)\,d\beta. \tag{5}$$

$Pr(\alpha)$ is thus the projection of $\Lambda(\alpha, \beta)$ onto the α axis.

The 1D function $Pr(\alpha)$ can be approximated by the discrete values $Pr(\alpha_i)$, where the α_i are the mid-points of a uniform grid with steps $\Delta\alpha$. In contrast, because of linearity, $\mathcal{N}_{i\ell}$ values $\beta_{i\ell}$ can readily be derived that randomly sample $Pr(\beta|\alpha_i)$. Combining these approaches, we derive the following approximation for the posterior distribution:

$$\Lambda(\alpha, \beta) = \sum_{i\ell} \Delta\alpha\,Pr(\alpha_i) \times \mathcal{N}_{i\ell}^{-1}\,\delta(\beta - \beta_{i\ell}). \tag{6}$$

With this approximation, all the quantities we wish to infer from the posterior distribution become weighted summations over the points $(\alpha_i, \beta_{i\ell})$, and these summations converge to exact values as $\Delta\alpha \rightarrow 0$ and $\mathcal{N}_{i\ell} \rightarrow \infty$. Arbitrary accuracy can therefore be achieved.

3.2. Astrometry only

If we only have astrometric data, Λ is a function of seven parameters. With the Thiele-Innes parameterization, the parameter vector is (ϕ, ψ), and the mathematical model is linear in the four ψ parameters and non-linear in the three ϕ parameters.

Following the 2D example of Sect. 3.1, we apply the chain rule to obtain

$$\Lambda(\phi, \psi|D_a) = Pr(\phi|D_a)\,Pr(\psi|\phi, D_a), \tag{7}$$

where

$$Pr(\phi|D_a) = \int \Lambda\,d\psi. \tag{8}$$

Here $Pr(\phi|D_a)$ is the projection of the 7D posterior distribution $\Lambda(\phi, \psi|D_a)$ onto the 3D ϕ-space. The second factor $Pr(\psi|\phi, D_a)$ then specifies how this projected or summed probability is to be distributed in ψ-space.

3.3. Astrometry and spectroscopy

With spectroscopic data included, Λ is now a function of 10 parameters (ϕ, ψ, λ). Again applying the chain rule, we write

$$\begin{aligned} \Lambda(\phi, \psi, \lambda|D_a, D_s) = Pr(\phi|D_a, D_s) \\ \times Pr(\psi|\phi, D_a, D_s)\,\times Pr(\lambda|\phi, \psi, D_s) \end{aligned} \tag{9}$$

where

$$Pr(\phi|D_a, D_s) = \int \Lambda\,d\psi\,d\lambda, \tag{10}$$

and

$$Pr(\psi|\phi, D_a, D_s) = \int \Lambda\,d\lambda \left/ \int\int \Lambda\,d\psi\,d\lambda \right. . \tag{11}$$

Here $Pr(\phi|D_a, D_s)$ is the projection of the 10D posterior distribution $\Lambda(\phi, \psi, \lambda|D_a, D_s)$ onto the 3D ϕ-space. The product $Pr(\psi|\phi, D_a, D_s) \times Pr(\lambda|\phi, \psi, D_s)$ then specifies how this summed probability is to be distributed first into ψ-space and then into λ-space.

The dependence of these probability factors on D_a and D_s merits comment. Both data sets contain information on $\phi = (\log P, e, \tau)$. Accordingly, $Pr(\phi)$ depends on both D_a and D_s.

The ψ-vector (A, B, F, G) determines the Campbell elements (a, i, ω, Ω) and vice versa. Since ω is a spectroscopic as well as an astrometric element, $Pr(\psi|\phi)$ must depend on D_s as well as on D_a.

If ϕ and ψ are given, then, since $\omega = \omega(\psi)$, the spectroscopic elements, P, e, τ, ω are known. The data D_s then suffices to determines the remaining spectroscopic elements $\lambda = (\gamma, K_1, K_2)$. Thus $Pr(\lambda|\phi, \psi)$ does not depend on D_a.

4. Likelihoods

The probability factors defined in Sect. 3 are now evaluated using Bayes' theorem and the appropriate likelihoods. Throughout this paper, we assume weak, non-informative priors whose impact on posterior distributions can be neglected.

4.1. Astrometry only

In this case, the posterior distribution is

$$\Lambda(\phi, \psi|D_a) \propto \mathcal{L}_a, \tag{12}$$

where, ignoring a constant factor,

$$\mathcal{L}_a = \exp\left(-\frac{1}{2}\chi_a^2\right), \tag{13}$$

and

$$\chi_a^2 = \Sigma_n w_n^a (x_n - \widetilde{x}_n)^2 + \Sigma_n w_n^a (y_n - \widetilde{y}_n)^2. \tag{14}$$

Because of linearity, $\widehat{\psi}(\phi)$, the minimum-χ^2 Thiele-Innes vector at given ϕ, is obtained without iteration, and we can write

$$\chi_a^2(\psi|\phi) = \widehat{\chi}_a^2(\widehat{\psi}|\phi) + \delta\chi_a^2(\delta\psi|\phi), \tag{15}$$

where $\delta\chi_a^2$ is the positive increment in χ_a^2 due to the displacement to $\psi = \widehat{\psi} + \delta\psi$.

Correspondingly, we write

$$\mathcal{L}_a(\phi, \psi) = \widehat{\mathcal{L}}_a(\phi)\,\widetilde{\mathcal{L}}_a(\psi|\phi), \tag{16}$$

where

$$\widehat{\mathcal{L}}_a = \exp\left(-\frac{1}{2}\widehat{\chi}_a^2\right) \quad \text{and} \quad \widetilde{\mathcal{L}}_a = \exp\left(-\frac{1}{2}\delta\chi_a^2\right). \tag{17}$$

The statistics of displacements in ψ-space is treated in Appendix A of L14b. These follow a quadrivariate normal distribution such that

$$Pr(\psi|\phi, D_a) = C^{-1}\exp\left(-\frac{1}{2}\delta\chi_a^2\right), \tag{18}$$

where $C(\phi) = (2\pi)^2\sqrt{\Delta}$ and Δ is the determinant of the covariance matrix. It follows that

$$\widetilde{\mathcal{L}}_a = C(\phi)Pr(\psi|\phi, D_a). \tag{19}$$

Substituting $\Lambda \propto \widehat{\mathcal{L}}_a\widetilde{\mathcal{L}}_a$ into Eq. (8) and eliminating $\widetilde{\mathcal{L}}_a$ with Eq. (19), we obtain

$$Pr(\phi|D_a) \propto C(\phi)\exp\left(-\frac{1}{2}\widehat{\chi}_a^2\right). \tag{20}$$

This determines the relative weights of the grid points ϕ_{ijk} and agrees with Eq. (14) in L14b.

From a random sampling of the quadrivariate normal distribution $Pr(\psi|\phi, D_a)$ we obtain the approximation

$$Pr(\psi|\phi, D_a) = \mathcal{N}_\psi^{-1}\sum_\ell \delta(\psi - \psi_\ell). \tag{21}$$

Accordingly, the relative weights from Eq. (20) are distributed equally among the points ψ_ℓ in ψ-space (note that at each ϕ_{ijk} an independent sample $\{\psi_\ell\}$ is drawn).

If the errors in $\widetilde{x}_n$ and $\widetilde{y}_n$ are uncorrelated, the quadrivariate distribution $Pr(\psi|\phi, D_a)$ simplifies to the product of two bivariate normal distributions – see Appendix A in L14b.

4.2. Astrometry and spectroscopy

With the addition of spectroscopic data and again assuming non-informative priors, the posterior density is

$$\Lambda(\phi, \psi, \lambda|D_a, D_s) \propto \mathcal{L}_a\mathcal{L}_s, \tag{22}$$

where, ignoring a constant factor,

$$\mathcal{L}_s = \exp\left(-\frac{1}{2}\chi_s^2\right), \tag{23}$$

and

$$\chi_s^2 = \Sigma_n w_{1n}^s (v_{1n} - \widetilde{v}_{1n})^2 + \Sigma_n w_{2n}^s (v_{2n} - \widetilde{v}_{2n})^2. \tag{24}$$

Because of linearity in $\lambda = (\gamma, K_1, K_2)$ when ϕ and ψ are fixed, $\widehat{\lambda}(\phi, \psi|D_s)$, the minimum-$\chi^2$ vector, is obtained without iteration, and we can write

$$\chi_s^2(\lambda|\phi, \psi) = \widehat{\chi}_s^2(\widehat{\lambda}|\phi, \psi) + \delta\chi_s^2(\delta\lambda|\phi, \psi), \tag{25}$$

where $\delta\chi_s^2$ is the positive increment in χ_s^2 due to the displacement to $\lambda = \widehat{\lambda} + \delta\lambda$.

Correspondingly, we write

$$\mathcal{L}_s(\lambda|\phi, \psi) = \widehat{\mathcal{L}}_s(\phi, \psi)\,\widetilde{\mathcal{L}}_s(\lambda|\phi, \psi), \tag{26}$$

where

$$\widehat{\mathcal{L}}_s = \exp\left(-\frac{1}{2}\widehat{\chi}_s^2\right) \quad \text{and} \quad \widetilde{\mathcal{L}}_s = \exp\left(-\frac{1}{2}\delta\chi_s^2\right). \tag{27}$$

The statistics of displacements in λ-space is treated in Appendix A. These follow a trivariate normal distribution $Pr(\lambda|\phi, \psi, D_s)$ such that Eq. (A.1) holds. From Eqs. (27) and (A.1), we obtain

$$\widetilde{\mathcal{L}}_s = \mathcal{D}(\phi, \psi)Pr(\lambda|\phi, \psi, D_s). \tag{28}$$

The statistics of displacements in ψ-space is modified by the spectroscopic data as noted in Sect. 3.3, and this is treated in Appendix B.

We now calculate $Pr(\phi|D_a, D_s)$. Substituting $\Lambda \propto \widehat{\mathcal{L}}_a\widetilde{\mathcal{L}}_a\widehat{\mathcal{L}}_s\widetilde{\mathcal{L}}_s$ into Eq. (10), eliminating $\widetilde{\mathcal{L}}_s$ with Eq. (28), and intergrating over λ, we obtain

$$Pr(\phi|D_a, D_s) \propto \widehat{\mathcal{L}}_a\int \mathcal{D}\,\widetilde{\mathcal{L}}_a\widehat{\mathcal{L}}_s\,d\psi. \tag{29}$$

We now eliminate $\widetilde{\mathcal{L}}_a$ using Eq. (19) to obtain

$$Pr(\phi|D_a, D_s) \propto C\widehat{\mathcal{L}}_a \int \mathcal{D}\,\widehat{\mathcal{L}}_s Pr(\psi|\phi, D_a)\,\mathrm{d}\psi. \qquad (30)$$

If we now replace $Pr(\psi|\phi, D_a)$ by the approximation given in Eq. (21) and assume that $\mathcal{N}_\psi$ is independent of ϕ, then

$$Pr(\phi|D_a, D_s) \propto C\widehat{\mathcal{L}}_a \sum_\ell (\mathcal{D}\widehat{\mathcal{L}}_s)_{\psi_\ell}. \qquad (31)$$

Accordingly, the relative weights of points ϕ_{ijk} in the ϕ-grid are

$$\mu_{ijk} = C\widehat{\mathcal{L}}_a \times \sum_\ell (\mathcal{D}\widehat{\mathcal{L}}_s)_{\psi_\ell}. \qquad (32)$$

Here the first factor $C\widehat{\mathcal{L}}_a$ depends only on D_a. The dependence on D_s is introduced by the second factor: If at a given ϕ, all ψ_ℓ correspond to poor fits to D_s, then the second factor disfavours that ϕ.

5. Numerical results

The technique developed in Sects. 3 and 4 is now applied to synthetic data D_a and D_s created as described in Sect. 2.3 for the model binary defined in Sect. 2.2. All calculations use a 100^3 grid for ϕ-space, and Monte Carlo sampling with $\mathcal{N}_\psi = 20$ for ψ-space and $\mathcal{N}_\lambda = 20$ for λ-space.

5.1. Parameter cloud

Let ϕ_{ijk} denote cell mid-points of the 3D grid spanning ϕ-space. At each ϕ_{ijk}, the technique generates $\mathcal{N}_\psi$ points ψ_ℓ in ψ-space. Then, at each ψ_ℓ, the technique generates $\mathcal{N}_\lambda$ points λ_m in λ-space. Thus, with this cascade, a cloud of points $(\phi_{ijk}, \psi_\ell, \lambda_m)$ is generated in the 10D (ϕ, ψ, λ)-space.

Note that this is a cloud of orbit parameters and *not* a cloud of orbits. Because linearity in ψ and λ is fully exploited, the values of χ_a^2 and χ_s^2 at cloud points are derived without computing astrometric and spectroscopic orbits, and this is the origin of the technique's computational efficiency.

The relative weights of cloud points $(\phi_{ijk}, \psi_\ell, \lambda_m)$ are

$$\mu_{ijk,\ell,m} = \mu_{ijk} \times \zeta_\ell \times \mathcal{N}_\lambda^{-1}. \qquad (33)$$

The first factor comes from Eq. (32), the second from Eq. (B.5), and the third from Eq. (A.4).

Note that the third factor is only relevant if $\mathcal{N}_\lambda$ varies with (ϕ, ψ). Note also that if $Pr(\psi|\phi, D_a, D_s)$ given by Eq. (B.2) were randomly sampled, the second factor would be $\mathcal{N}_\psi^{-1}$. Instead, the quadrivariate normal distribution $Pr(\psi|\phi, D_a)$ is sampled and then corrected via the coefficients ζ_ℓ, which are such that $\sum \zeta_\ell = 1$ – see Appendix B.

5.2. Inferences

Suppose $Q(\Theta)$ is a quantity of interest. Its posterior distribution derived from the parameter cloud is

$$\Theta(Q) = \sum_t \mu_t\,\delta(Q - Q_t) \bigg/ \sum_t \mu_t\,, \qquad (34)$$

where $t \equiv (ijk, \ell, m)$. The corresponding cumulative distribution function is
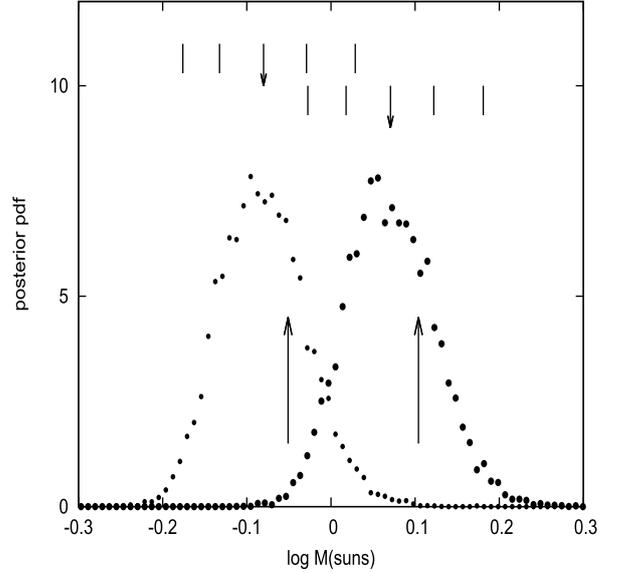
Fig. 1. Posterior densities for $\log \mathcal{M}_1$ and $\log \mathcal{M}_2$. The long vertical arrows indicate exact values. The short vertical arrows and lines indicate the posterior means and the 1- and 2-σ credibility intervals.

$$F(Q) = \sum_{Q_t < Q} \mu_t \bigg/ \sum_t \mu_t\,. \qquad (35)$$

The *equal tail* credibility interval (Q_L, Q_U) corresponding to $\pm 1\sigma$ is then obtained from the equations

$$F(Q_L) = 0.1587 \qquad F(Q_U) = 0.8413, \qquad (36)$$

so that the enclosed probability is 0.6826.

These credibility intervals are asymptotically rigorous – i.e., are exact in the limits $\mathcal{N}_\psi \to \infty$, $\mathcal{N}_\lambda \to \infty$, and grid steps $\to 0$.

5.3. An example

The fundamental data derivable from the combined astrometric and spectroscopic data are the component masses $\mathcal{M}_1, \mathcal{M}_2$ and the parallax ϖ. None of these quantities can be derived if only one data set is available.

At every cloud point t, Kepler's law and the two spectroscopic mass functions can be solved for $\mathcal{M}_1, \mathcal{M}_2$ and ϖ, and these values each have relative weight μ_t. Accordingly, the posterior densities of these quantities can be calculated from Eq. (34) and their credibility intervals from Eq. (35).

For a particular simulation of D_a and of D_s, the posterior densities so derived are plotted in Figs. 1 and 2. Also plotted are the posterior means, the 1– and 2– σ credibility intervals, as well as the exact values from Sect. 2.2. In each case, the exact values fall within the 1– σ limits.

In Appendix C, following L16, a Bayesian goodness-of-fit (GOF) statistic, χ_B^2, is defined together with corresponding Bayesian p-value. We now apply this test. For the astrometric data, the posterior mean of χ_a^2 is $\langle \chi_a^2 \rangle_u = 22.2$, and for the spectroscopic data $\langle \chi_s^2 \rangle_u = 16.9$, so that in total $\langle \chi^2 \rangle_u = 39.1$. Since the total number of parameters is $k = 10$, Eq. (C.4) gives $\chi_B^2 = 29.1$.

The total number of measurements is $n = 40$, comprising two astrometric (x, y) and two spectroscopic (v_1, v_2) measurements in each of ten years. The number of degrees of freedom is therefore $\nu = n - k = 30$. Substitution in Eq. (C.5) then gives $p_B = 0.51$, a value consistent with the belief that the data is analysed with a valid model and that Bayesian inferences are not suspect.
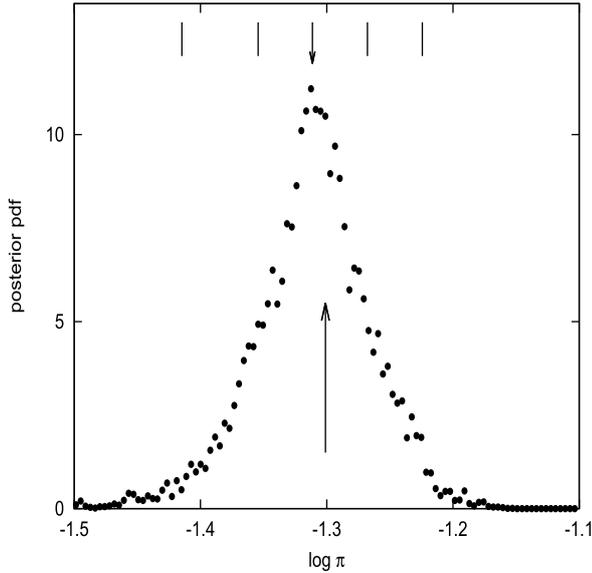
Fig. 2. Posterior density for $\log \varpi('')$. The long vertical arrow indicates the exact value. The short vertical arrow and lines indicate the posterior mean and the 1- and 2-σ credibility intervals.

Table 1. Coverage fractions from 10^3 trials.

Q	$1 - \sigma$	$2 - \sigma$
$\mathcal{E}(f)$	0.683 ± 0.015	0.954 ± 0.007
$\log \mathcal{M}_1$	0.711 ± 0.014	0.964 ± 0.006
$\log \mathcal{M}_2$	0.732 ± 0.014	0.963 ± 0.006
$\log \varpi$	0.667 ± 0.015	0.963 ± 0.006

5.4. Coverage

An accurate gauge of the statistical performance of the technique requires many repetitions of the above calculation with independently drawn samples of D_a and D_s.

With different data, the posterior densities and corresponding credibility limits in Figs. 1 and 2 change. But the long vertical arrows marking exact values remain fixed. For each independent repetition, we can record whether or not the exact values are enclosed by the 1- and 2-σ credibility intervals. In this way, we carry out a coverage experiment as in L14a,b – see also Martinez et al. 2017.

The results obtained from 1000 repetitions are summarized in Table 1. These show reasonable agreement with $\varepsilon(f)$, the expected fractions for errors obeying a normal distribution. Thus, despite the non-linearities, the credibility intervals retain their conventional interpretations.

6. Discussion: Bayesian hypothesis testing

In L16 and references therein, the relative absence in the astronomical literature of statistical testing of Bayesian models is commented upon.

6.1. Some quotes

The following quote from a statistician (Anscombe 1963) indicates that concern on this issue is of long standing:

"To anyone sympathetic with the current neo-Bernoullian neo-Bayesian Ramseyesque Finettist Savageous movement in statistics, the subject of testing goodness-of-fit is something of an embarrassment."

A very recent comment (Fischer et al. 2016 in Sect. 4.1, authored by E. Ford) referring to exoplanets is:

"Too often people using Bayesian methods ignore model checking, because it doesn't have a neat and tidy formal expression in the Bayesian approach. But it is no less necessary to do GOF type checks for a Bayesian analysis than it is for a frequentist analysis"

6.2. Additional justifications

When authors ignore model checking, they seldom, if ever, explain why. The quote above suggests that Bayesians are deterred by the absence of a readily-applied test. In contrast, frequentists reporting a minimum-χ^2 analysis generally include χ_0^2, the χ^2 minimum, and often also the p-value derived from the known distribution of χ_0^2. Thus, this traditional, frequentist approach has a *built-in* reality check. Moreover, this check is rigorously justified for linear models and normally-distributed measurement errors.

Note that minimum-χ^2 codes return estimates and confidence intervals even when χ_0^2 corresponds to a vanishingly small p-value. Thus, we may surmise that innumerable spurious inferences from false hypotheses or poor data are absent from the scientific literature precisely because of this built-in check.

Besides the difficulty of Bayesian model-checking, it seems likely that the following additional reasons play a role in checking being ignored: The detection of the expected signal confirms the hypothesis.

This is endemic in studies of orbits, including frequentist analyses going back decades. If a star is investigated for reflex motion due to a companion and a periodic signal is detected, then it is all too easy to take this as confirmation of a companion. A more critical approach recognizes that a harmonic expansion of Keplerian motion provides quantitative tests of the orbit hypothesis. When this approach is applied to a sample of spectroscopic binaries with exquisitely accurate radial velocities, significant departures from exact Keplerian motion are found (Lucy 2005; Hearnshaw et al. 2012).

A notable recent signal detection is that of gravitational waves from coalescing black holes (Abbott et al. 2016). The published parameters for the initial black hole binary derive from a Bayesian analysis. But these authors do not ignore model checking: their Bayesian analysis is preceded by a standard frequentist χ^2 test of template fits. The Bayesian model has so many parameters that a poor fit is improbable.

In this case, the acceptance-rejection aspect of the scientific method is replaced by the posterior density favouring or disfavouring regions of parameter space. The expectation is that with enough high quality data, the posterior density will be sharply peaked at the point corresponding to the true hypothesis. But what if the true hypothesis is not part of the adopted multi-parameter model? How does the investigator detect this? A hypothesis should not be rejected if there is no alternative.

On this view, Bayesian model checking, even if readily carried out, should not lead to the rejection of a hypothesis. Rather, one should wait until an alternative hypothesis is proposed and then implement the model selection machinery. This view goes back to (Jeffreys 1939, Sect. 7.2.2) – see also (Sivia & Skilling 2006, p.85).

Jeffreys supports this view by remarking that there was never a time over previous centuries when Newton's theory of gravity would not have failed a p-test. It is therefore instructive to recall

how Adams and Le Verrier reacted to the large residuals in the motion of Uranus – i.e., small *p*-value. Crucially, their view was that the hypothesis being tested was not Newton's theory but the then current 7-planet model of the solar system. Because they had a far greater degree of belief in Newtonian gravity than in the 7-planet model, they doubted the latter and went on to successfully predict Neptune's position. This example illustrates that ambitious and effective scientists take small *p*-values seriously even in the absence of alternative hypotheses. By doing so, they create alternative hypotheses.

6.3. The χ_B^2 statistic

A Bayesian GOF statistic χ_B^2 and corresponding *p*-value is defined in Appendix C.

A crucial requirement of any (GOF) statistic is that it should not often falsely lead one to reject or doubt an hypothesis when that hypothesis is true. In frequentist terms, this is a Type II error. Such an error arises when the statistic gives a small *p*-value, say $p = 0.001$, even when the null hypothesis (H_0) is true. Of course, such a value can occur by chance, but for an acceptable GOF statistic the frequency of Type II errors should not markedly exceed *p*.

In the simulation reported in Sects. 5.3 and 5.4, the null hypothesis is correct by construction, since the data is generated from the exact formulae for the astrometric and spectroscopic orbits. Thus, if the mathematical models were completely linear, we would expect χ_B^2 to be distributed exactly as χ_ν^2 with $\nu = n - k$ degrees of freedom. The *p*-value defined by Eq. (C.5) would then have an exactly uniform distribution in the interval $(0, 1)$.

The $N_{tot} = 1000$ simulations used for the coverage experiment in Sect. 5.4 allow the uniformity of the p_B values to be tested. In Fig. 3, the fraction with $p_B < p$ is plotted against *p*. We see that uniformity is obeyed with reasonable precision for $p \in (0.01, 1.00)$. In particular, there is no indication of a significant departure from uniformity that could be attributed to the non-linearities. Accordingly, the p_B values derived from χ_B^2 can be interpreted in the same way and with the same confidence as *p*-values in minimum-χ^2 estimation.

Note that the calculation of χ_B^2 is a trivial addition to an existing Bayesian code that very likely already calculates the posterior means of other quantities.

6.4. Posterior predictive *p*-values

In the contribution authored by E. Ford from which the quote in Sect. 6.1 is taken, readers are referred to Gelman et al. (2013) who recommend posterior predictive *p*-values for testing Bayesian models.

In the context of the technique developed here, this recommendation proceeds as follows:
1) Randomly select a point in parameter space from the posterior distribution. Thus, if *t* is an index that gives a 1D enumeration of the parameter cloud, a random point t' is that which most closely satisfies the equation

$$\sum_{t < t'} \mu_t \Big/ \sum_t \mu_t = x, \tag{37}$$

where *x* is a random number $\in (0, 1)$.
2) From the 10 parameters at t', create synthetic data $D' = D'_a + D'_s$, compute $\chi^2 = \chi_a^2 + \chi_s^2$, and then compare to the χ^2 at t' for the original data $D = D_a + D_s$.
3) Repeat steps 1) and 2) N_{tot} times.
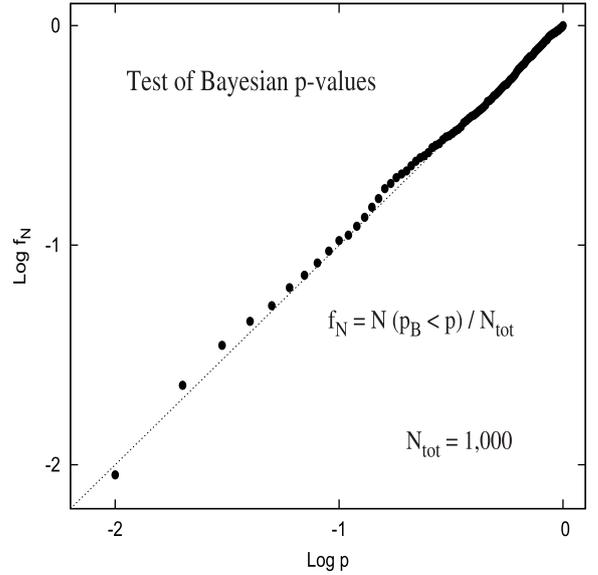
Fig. 3. Test of Bayesian *p*-values. From 1,000 simulations, the fraction with $p_B > p$ is plotted against *p* for $p = 0.01(0.01)1.00$. The dashed line shows the expected result if the null hypothesis is correct *and* if the statistic χ_B^2 obeys the χ_ν^2 distribution with $\nu = n - k$ degrees of freedom.

A Bayesian *p*-value is then defined to be

$$p_B = \mathcal{N}[\chi^2(D') > \chi^2(D)] / N_{tot}. \tag{38}$$

Thus a small value of p_B indicates that it is hard to find points t' giving a worse fit than the original data, indicating that original data gives a poor fit.

To quote E. Ford again (Sect. 6.1), posterior predictive checking is evidently not "a neat and tidy" formalism. Moreover, physical scientists have a strong interest in having reliable *p*-values when $p \lesssim 0.001$, since such values raise serious doubts about a model's validity. This then requires $\mathcal{N}_{tot} \sim 10\,000$ repetitions of the above steps, which may be infeasible.

Posterior predictive *p*-values have been compared to the values given by Eq. (C.5) for a simple 1D toy model. Specifically, a Hubble flow in Euclidean space populated by perfect standard candles. Synthetic data is created and the posterior density for the Hubble constant derived. A poor fit can be engineered by corrupting the data at high redshift and then comparing the resulting two small *p*-values. They agree closely.

This suggests that the readily calculated p_B given by Eq. (C.5) eliminates any need for the cumbersome direct calculation of the posterior predictive *p*-value given by Eq. (38).

7. Conclusion

In this paper, a non-trivial example of a wide class of problems in statistical astronomy is addressed. These are the so-called hybrid problems where the mathematical models predicting the observations are partly linear and partly non-linear in the basic parameters. As in the simpler, purely astrometric case considered in L14b, when spectroscopic data is added, a grid search over the non-linear parameter space combined with Monte Carlo sampling in the linear parameter spaces still leads to a computationally efficient scheme and again yields credibility intervals with close to correct coverage (Sect. 5.4), a result of prime importance generally, but especially so when estimating fundamental stellar parameters.

In contrast to L14a, the formulation in Sects. 3 and 4 is mostly quite general and so should be readily adapted to other hybrid problems.

In addition to exhibiting correct coverage, the large number of independent simulations allow the testing (Sect. 6.3) of χ_B^2, a Bayesian GOF criterion (Appendix C) for posterior probability densities. Even though the test problem involves some non-linear parameters, the exact sampling distribution in the linear case is closely followed, thus providing a readily-calculated p-value that quantifies one's confidence in the inferences drawn from the posterior distribution. Since in problems that are exactly linear, the Bayesian and frequentist p-values are identical, investigators can make the decisions on the basis of the Bayesian p_B-value exactly as they would for a frequentist p-value. Moreover, since calculating the statistic χ_B^2 involves trivial changes to a Bayesian code, it provides "the neat and tidy formal expression" that is missing in current Bayesian methodology – see quote from E. Ford in Sect. 6.1.

References

Abbott, B.P., et al. (LIGO Scientific & Virgo Collaborations) 2016, Phys. Rev. Lett., 116, 061102
Anscombe, F. J. 1963 J. Roy. Stat. Soc. Ser. B, 25, 81
Catanzarite, J. 2010, ArXiv e-prints [arXiv:1008.3416]
Eastman, J., Gaudi, B., & Agol, E. 2013, PASP, 125, 83
Fischer, D. A., Anglada-Escude, G., Arriagada, P., et al. 2016, PASP, 128, 6001
Gelman, A., Carlin, J. B., Stern, H. S., et al. 2013, Bayesian Data Analysis (3rd edn.), ed. Chapman and Hall (Boca Raton, FL: CRC Press)
Gentle, J. E. 2009, Computational Statistics (New York: Springer), 315
Hamilton, W.C. 1964, Statistics in Physical Science (New York: Roland Press)
Hearnshaw, J. B., Komonjinda, S., Skuljan, J. & Kilmartin, P.M. 2012, MNRAS, 427, 298
Jeffreys, H. 1939, Theory of Probability (Oxford: Clarendon Press)
Lucy, L. B. 2005, A&A, 439, 663
Lucy, L. B. 2014a, A&A, 563, A126 (L14a)
Lucy, L. B. 2014b, A&A, 565, A37 (L14b)
Lucy, L. B. 2016, A&A, 588, A19 (L16)
Martinez, G. D., Kosmo, K., Hees, A., Ahn, J., & Ghez, A. 2017, IAU Symp., 322, 239
Mason, B. D., Douglass, G.G. & Hartkopf, W.I. 1999, AJ, 117, 1023
Press W. H., Teukolsky S. A., Vetterling W. T., & Flannery B. P. 2007, Numerical Recipes 3rd edn. (Cambridge: Cambridge Univ. Press)
Sivia, D. S., & Skilling, J. 2006, Data Analysis, A Bayesian Tutorial 2nd edn., (Oxford University Press)
Wright, J. T., & Howard, A. W. 2009, ApJS, 182, 205

Appendix A: Statistics in λ-space

Statistics in the 4D Thiele-Innes ψ-space is treated in Appendix A of L14b. Analogous results are briefly stated here for the 3D λ-space.

Given ϕ and ψ, the minimum-χ_s^2 vector $\widehat{\lambda} = (\widehat{\gamma}, \widehat{K_1}, \widehat{K_2})$ is obtained without iteration from the normal equations derived from Eqs. (1) and (24).

The displacement $\lambda = \widehat{\lambda} + \delta\lambda$ gives $\chi_s^2 = \widehat{\chi_s^2} + \delta\chi_s^2$ with positive $\delta\chi_s^2$. On the assumption of normally-distributed errors, the probability density at λ is a trivariate normal distribution such that

$$Pr(\lambda|\phi, \psi, D_s) = \mathcal{D}^{-1} \exp(-\frac{1}{2}\delta\chi_s^2), \tag{A.1}$$

where $\mathcal{D}(\phi, \psi) = (2\pi)^{3/2} \sqrt{\Delta}$ and Δ is the determinant of the covariance matrix.

A.1. Random sampling in λ-space

Points λ_ℓ randomly sampling the trivariate normal distribution $Pr(\lambda|\phi, \psi, D_s)$ are derived with a standard procedure (Gentle 2009) for sampling multivariate distributions. The first step is to make a Cholesky decomposition (Press et al. 2007, p. 100) of the covariance matrix C – i.e., to find the lower triangular matrix L such that

$$LL' = C. \tag{A.2}$$

A random sample from $Pr(\lambda|\phi, \psi, D_s)$ is then

$$\lambda = \widehat{\lambda} + L.z = \widehat{\lambda} + \delta\lambda, \tag{A.3}$$

where the elements of $z = (z_1, z_2, z_3)$ are random Gaussian variates. The resulting approximation to $Pr(\lambda)$ is

$$Pr(\lambda|\phi, \psi, D_s) = \mathcal{N}_\lambda^{-1} \sum_m \delta(\lambda - \lambda_m). \tag{A.4}$$

Note that $\mathcal{N}_\lambda$ can vary with (ϕ, ψ). The increment in χ_s^2 due to the displacement from $\widehat{\lambda}$ is

$$\delta\chi_s^2 = z_1^2 + z_2^2 + z_3^2. \tag{A.5}$$

Accordingly, as in the analogous problem in ψ-space – see Eq. (A.22) in L14b, the increment in χ^2 is obtained without computing the spectroscopic orbits at $\widehat{\lambda} + \delta\lambda$ – though this should be checked during code development. This is a consequence of linearity and accounts for the computational efficiency of the technique.

In Appendix A of L14, Cholesky decompostion is not needed because the quadrivariate normal distribution $Pr(\psi|\phi, D_a)$ is the product of bivariate distributions. But this simplification is lost if $\widetilde{x}_n$ and $\widetilde{y}_n$ have correlated errors (Sect. 2.3). In that circumstance, the above Cholesky approach is the necessary generalization.

Appendix B: Modified statistics in ψ- space

The treatment of statistics in ψ-space in Appendix A of L14b does not apply when spectroscopic data is included. As noted in Sect. 3.3 – see Eq. (11), $Pr(\psi)$ depends on both D_a and D_s

The required modification is obtained by substituting $\Lambda \propto \widehat{\mathcal{L}_a}\widetilde{\mathcal{L}_a}\widehat{\mathcal{L}_s}\widetilde{\mathcal{L}_s}$ into Eq. (11), integrating over λ using Eq. (28), and noting that $\widehat{\mathcal{L}_a}$ is independent of ψ. This gives

$$Pr(\psi|\phi, D_a, D_s) \propto \mathcal{D} \widetilde{\mathcal{L}_a}\widehat{\mathcal{L}_s}. \tag{B.1}$$

We now eliminate $\widetilde{\mathcal{L}_a}$ using Eq. (19) and noting that C is independent of ψ. This gives

$$Pr(\psi|\phi, D_a, D_s) \propto \mathcal{D} \widehat{\mathcal{L}_s} Pr(\psi|\phi, D_a), \tag{B.2}$$

showing that $Pr(\psi)$ is modified from the pure astrometry case by the factor $\mathcal{D}\widehat{\mathcal{L}_s}$ introduced by spectroscopy. Because of this modification, $Pr(\psi|\phi, D_a, D_s)$ is not a multivariate normal distribution and so not as readily sampled.

The adopted sampling procedure is as follows: from Eq. (A.20) in L14b, we have the approximation

$$Pr(\psi|\phi, D_a) = \mathcal{N}_\psi^{-1} \sum_\ell \delta(\psi - \psi_\ell), \tag{B.3}$$

where the ψ_ℓ randomly sample the quadrivariate normal distribtion appropriate in the pure astrometry case (Appendix A, L14b). Substituting into Eq. (B.2), we obtain the corresponding approximation when spectroscopy is included

$$Pr(\psi|\phi, D_a, D_s) = \sum_\ell \zeta_\ell \, \delta(\psi - \psi_\ell), \tag{B.4}$$

where

$$\zeta_\ell(\phi) = (\mathcal{D}\widehat{\mathcal{L}_s})_{\psi_\ell} \Big/ \sum_\ell (\mathcal{D}\widehat{\mathcal{L}_s})_{\psi_\ell}. \tag{B.5}$$

As might be expected, because of the factor $(\widehat{\mathcal{L}_s})_{\psi_\ell}$, points ψ_ℓ in ψ-space are strongly disfavoured if that ψ_ℓ gives a poor fit to the spectroscopic data.

Appendix C: The χ_B^2 and ψ_B^2 statistics

In an earlier paper (Lucy 2016; L16), we define

$$\langle \chi^2 \rangle_u = \int \chi^2(\alpha) \Lambda(\alpha|D) \, \mathrm{d}V_\alpha, \tag{C.1}$$

where

$$\Lambda(\alpha|D) = \mathcal{L}(\alpha|D) \Big/ \int \mathcal{L}(\alpha|D) \, \mathrm{d}V_\alpha. \tag{C.2}$$

Thus $\langle \chi^2 \rangle_u$ is the posterior mean of $\chi^2(\alpha)$ when the posterior density Λ is computed under the assumption of a uniform (u) prior.

If the model is linear in the parameter vector α and if errors are normally-distributed, then (Appendix A, L16)

$$\langle \chi^2 \rangle_u = \chi_0^2 + k, \tag{C.3}$$

where χ_0^2 is the minimum value of $\chi^2(\alpha)$ and k is the number of parameters. Moreover, under the stated assumptions, χ_0^2 is distributed as χ_ν^2, where $\nu = n - k$ is the number of degrees of freedom and n is the number of measurements.

It follows that if we define the statistic

$$\chi_B^2 = \langle \chi^2 \rangle_u - k, \tag{C.4}$$

then, for a linear model and normally-distributed errors, χ_B^2 is distributed as χ_ν^2 with $\nu = n - k$. Accordingly, a p-value that quantifies the quality of the posterior distribution $\Lambda(\alpha|D)$ from which all Bayesian inferences are drawn is given by

$$p_B = Pr(\chi_\nu^2 > \chi_B^2) \quad for \quad \nu = n - k. \tag{C.5}$$

If the model is indeed linear in α, it follows from Eqs. (C.3) and (C.4) that $\chi_B^2 = \chi_0^2$, and so the frequentist and Bayesian p-values agree, a gratifying result.

If the model is non-linear in some parameters, then this GOF test should still be useful if the data is such that the fractional error of the non-linear parameters are small, for then a linearized model could be used.

In most Bayesian analyses in astronomy, the imposed priors are uninformative and so this analysis holds. In the rare cases where an informative prior π is imposed, perhaps from a previous experiment, the discussion in L16, Sect. 4.1 suggests that the criterion χ_B^2 with $\langle \chi^2 \rangle_\pi$ replacing $\langle \chi^2 \rangle_u$ will have closely similar characteristics.

C.1. Generalization

The above analysis assumes uncorrelated measurement errors. The inclusion of correlations is a straightforward application of the statistics of quadratic forms – see, e.g., Hamilton (1964, Chap.4).

When corellations are included, the χ^2 summation is replaced by

$$\psi^2 = \boldsymbol{v}' \boldsymbol{M}^{-1} \boldsymbol{v}, \tag{C.6}$$

where $\boldsymbol{M}$ is the covariance matrix and $\boldsymbol{v}$ is the vector of residuals. The previous analysis assumes that the off-diagonal elements of $\boldsymbol{M}^{-1}$ are zero.

With linearity in the parameter vector $\boldsymbol{\alpha}$ and normally-distributed measurement errors $\tilde{x} - x$, the sampling distribution of $\tilde{x}$ is a k-dimensional multivariate normal disribution $\propto \exp(-\psi^2/2)$, where k is the number of parameters. It follows that the likelihood is

$$\mathcal{L}(\boldsymbol{\alpha}|D) \propto \exp(-\frac{1}{2}\psi^2). \tag{C.7}$$

Exploiting linearity in $\boldsymbol{\alpha}$ and assuming a weak prior, we can write the posterior density as

$$\Lambda(\boldsymbol{\alpha}|D) \propto \exp(-\frac{1}{2}\psi_0^2) \times \exp(-\frac{1}{2}\delta\psi^2), \tag{C.8}$$

where ψ_0^2 is the minimum of ψ^2 at α_0 and $\delta\psi^2$ is the positive increment due to the displacement $\boldsymbol{\alpha} - \boldsymbol{\alpha}_0$.

Accordingly, in the case of a uniform prior, the posterior mean of ψ^2 is

$$\langle \psi^2 \rangle_u = \psi_0^2 + \frac{\int \Delta\psi^2 \exp(-\Delta\psi^2/2)\, dV_\alpha}{\int \exp(-\Delta\psi^2/2)\, dV_\alpha}. \tag{C.9}$$

This has the same form as Eq. (A.4) in L16, with $\Delta\psi^2$ replacing $\Delta\chi^2$. Therefore, since surfaces of constant $\Delta\psi^2$ are also self-similar k-dimensional ellipsoids, we immediately have

$$\langle \psi^2 \rangle_u = \psi_0^2 + k. \tag{C.10}$$

Now ψ_0^2 is distributed as χ_ν^2 with $\nu = n - k$ degrees of freedom. Accordingly, the statistic

$$\psi_B^2 = \langle \psi^2 \rangle_u - k, \tag{C.11}$$

is distributed as χ_ν^2 with $\nu = n - k$ degrees of freedom.

This is the required generalization of χ_B^2 given by Eq. (C.4).