

Frequentist confidence intervals for orbits[★]

L. B. Lucy

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

Received 17 February 2014 / Accepted 25 March 2014

ABSTRACT

The problem of efficiently computing the orbital elements of a visual binary while still deriving confidence intervals with frequentist properties is treated. When formulated in terms of the Thiele-Innes elements, the known distribution of probability in Thiele-Innes space allows efficient grid-search plus Monte-Carlo-sampling schemes to be constructed for both the minimum- χ^2 and the Bayesian approaches to parameter estimation. Numerical experiments with 10^4 independent realizations of an observed orbit confirm that the 1- and 2σ confidence and credibility intervals have coverage fractions close to their frequentist values.

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

1. Introduction

When error bars or confidence intervals are reported, the reader expects them to have their frequentist meaning. Thus, a 95% confidence interval is interpreted as implying a probability of 0.95 that the true result is enclosed by that interval. Similarly, the interval defined by $\pm 1\sigma$ error bars is expected to include the true answer with probability 0.683. However, this frequentist ideal is often not realized. This may be the result of observers misjudging the precision of their measurements or of large measurement errors occurring more frequently than expected for a normal distribution.

Such practical issues are absent when data analysis techniques are investigated with simulations, since precision can be exactly specified and measurement errors can be assigned with random Gaussian variates, so that one might then expect a rigorous recovery of the frequentist ideal. But approximations can still compromise statistical rigour. For example, if a grid is required, confidence intervals might be affected if the grid is too coarse. In such cases, with increased computational resources, the limit as the grid steps $\rightarrow 0$ can be closely approached and accurate results obtained.

Of more concern are approximations that compromise confidence intervals independently of any such limit. Two examples in the recent literature occur in *hybrid* problems – i.e., non-linear problems with a subset of linear parameters. The first example is the code EXOFAST for analysing transit and radial velocity data for stars with orbiting planets (Eastman et al. 2013). These authors note that the convergence of their Markov chain Monte Carlo (MCMC) parameter search is much faster if the exact solution for the linear parameters is introduced. However, the resulting uncertainties in the linear parameters are as much as 10 times smaller than when fitted non-linearly. Pending further research, these authors sensibly choose the inefficient option of treating all parameters as non-linear.

A similar but less extreme example arises when Bayesian estimation is applied to visual binaries (Lucy 2014; L14). When formulated in terms of the Thiele-Innes elements, the problem

becomes linear in four of the seven elements. But when this linearity is exploited, coverage fractions (L14, Sect. 5.5) indicate that the standard errors of the four linear elements are too small by factors of up to 2.1.

These examples pose a statistical challenge in the analysis of orbits: How can we benefit from partial linearity without losing the frequentist properties of confidence intervals? In this paper, this challenge is addressed in its visual binary context and for both frequentist and Bayesian procedures.

2. Synthetic orbits

The paper L14 is followed closely with regard both to notation and in the creation of synthetic data.

2.1. Orbital elements

The orbit of the secondary relative to its primary is conventionally parameterized by its 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 L14 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 $\in (0, 1)$.)

2.2. Model binary

As in L14, the adopted model binary has the following Campbell elements:

$$\begin{aligned} P_* &= 100y & \tau_* &= 0.4 & e_* &= 0.5 & a_* &= 1'' \\ i_* &= 60^\circ & \omega_* &= 250^\circ & \Omega_* &= 120^\circ. \end{aligned} \quad (1)$$

An observing campaign for this binary is simulated by creating measured Cartesian sky coordinates $(\tilde{x}_n, \tilde{y}_n)$ with weights w_n at

[★] Appendix is available in electronic form at <http://www.aanda.org>

uniformly-spaced times t_n for $n = 1, \dots, N$ as described in L14 Sect. 3.2. The parameters defining a campaign are f_{orb} , the fraction of the orbit observed, N , the number of observations, and σ , the standard error for unit weight.

For given orbital elements, the predicted orbit (x_n, y_n) is computed as described in L14 Sect. A.1, and the quality of the fit is determined by

$$\chi^2 = \frac{1}{\sigma^2} \sum_n w_n (x_n - \tilde{x}_n)^2 + \frac{1}{\sigma^2} \sum_n w_n (y_n - \tilde{y}_n)^2. \quad (2)$$

3. Minimum- χ^2 estimation

The conventional (frequentist) approach to orbit-fitting is the method of least squares – i.e., finding the elements $\hat{\theta} = (\hat{\phi}, \hat{\psi})$ that minimize χ^2 . When the problem is non-linear, the search for the minimum typically involves successive differential corrections obtained from linearized equations, starting with an initial guess. However, in treating incomplete orbits and imprecise data, it is preferable to find $\chi_{\text{min}}^2 = \chi^2(\hat{\theta})$ by means of a grid search (e.g., Hartkopf et al. 1989; Schaefer et al. 2006) and then to derive confidence intervals from constant χ^2 “surfaces” in parameter space (e.g., Press et al. 1992, Chap. 15.6; James 2006, Chap. 9.1.2).

3.1. Grid search

In a brute force approach to finding $\hat{\theta}$, values of χ^2 would be computed throughout a 7D grid. Confidence intervals for the elements would then be derived from projections of the 7D volume \mathcal{V} defined by the inequality

$$\chi^2(\theta) < \chi_{\text{min}}^2 + \Delta_{\nu, \alpha} \quad (3)$$

where the constant $\Delta_{\nu, \alpha}$ is determined by ν , the number of degrees of freedom and α , the desired confidence level. With a typical 100 steps for each dimension, the brute force method requires χ^2 to be evaluated at $\sim 10^{14}$ grid points. However, if the linearity with respect to the Thiele-Innes elements can be exploited, χ^2 is only required at $\sim 10^6$ grid points. This potential reduction by a factor of $\sim 10^8$ in the number of computed orbits is a powerful incentive to solve the challenge posed in Sect. 1.

On the assumption that linearity can be exploited, grid searches in this paper are restricted to the ϕ -elements P, e, τ . The grid is defined by taking constant steps spanning the intervals $(\log P_L, \log P_U), (e_L, e_U), (\tau_L, \tau_U)$. The grid cells are labelled (i, j, k) and the ϕ -elements at the mid-points are $\log P_i, e_j, \tau_k$. With these values fixed, χ_{ijk}^2 is a function of $\psi = (A, B, F, G)$ and has its minimum value $\hat{\chi}_{ijk}^2$ at the point $\hat{\psi}_{ijk} = (\hat{A}, \hat{B}, \hat{F}, \hat{G})$ given by Eqs. (A.7) in L14.

Because $\chi_{ijk}^2 \geq \hat{\chi}_{ijk}^2$, it follows that nowhere in (ϕ_{ijk}, ψ) -space is $\chi_{ijk}^2(\psi)$ less than the minimum value found in the 3D search. Accordingly, in the limit of vanishingly small grid steps,

$$\chi_{\text{min}}^2 = \min_{\theta} \{\chi^2(\theta)\} = \min_{ijk} \{\hat{\chi}_{ijk}^2\}. \quad (4)$$

Thus, as has long been understood (e.g. Hartkopf et al. 1989), the minimum- χ^2 elements $\hat{\theta}$ can be found with a grid search restricted to the non-linear elements.

3.2. Approximate confidence intervals

With $\hat{\theta}$ determined, the calculation of confidence intervals requires projections of \mathcal{V} , the volume in θ -space defined by

Eq. (3). In the absence of a 7D grid, a possible approach is to derive approximate confidence intervals from projections of the 3D grid points satisfying Eq. (3) – i.e., from projections of the domain \mathcal{D} comprising grid points such that

$$\hat{\chi}_{ijk}^2 < \chi_{\text{min}}^2 + \Delta_{\nu, \alpha}. \quad (5)$$

This derivation of confidence intervals has exploited linearity since it relies on obtaining $\hat{\psi}_{ijk}$ and therefore also $\hat{\chi}_{ijk}^2$ without iteration. However, every point $\in \mathcal{D}$ also satisfies Eq. (3), so that $\mathcal{D} \in \mathcal{V}$. Accordingly, these approximate intervals will always be enclosed within the true intervals and so may give a misleading impression of the elements’ precision.

3.3. Accurate confidence intervals

An asymptotically rigorous calculation proceeds as follows: first, since $\chi_{ijk}^2(\psi) \geq \hat{\chi}_{ijk}^2$, the points (ϕ_{ijk}, ψ) are all exterior to \mathcal{V} when $\phi_{ijk} \ni \mathcal{D}$. Thus grid points $\ni \mathcal{D}$ are no longer of interest.

Now consider a point $\phi_{ijk} \in \mathcal{D}$. A point in the ψ -space attached to ϕ_{ijk} has a χ^2 higher than the least squares value at ϕ_{ijk} by the amount

$$\delta\chi_{ijk}^2 = \chi_{ijk}^2 - \hat{\chi}_{ijk}^2. \quad (6)$$

This point is on the 6D surface \mathcal{S} of the volume \mathcal{V} if

$$\delta\chi_{ijk}^2 = \Delta_{\nu, \alpha} - (\hat{\chi}_{ijk}^2 - \chi_{\text{min}}^2). \quad (7)$$

The contribution to \mathcal{S} arising from the grid point (i, j, k) can therefore be obtained by randomly sampling the attached ψ -space subject to this constraint on $\delta\chi_{ijk}^2$. The superposition of these contributions from all $\phi_{ijk} \in \mathcal{D}$ then maps out \mathcal{S} , and the projections of \mathcal{S} give the desired confidence intervals.

If \mathcal{N} is the number of random points ψ_ℓ on \mathcal{S} generated at each ϕ_{ijk} , the ensemble of generated points (ϕ_{ijk}, ψ_ℓ) becomes an exact representation of \mathcal{S} in the limits $\mathcal{N} \rightarrow \infty$ and grid steps $\rightarrow 0$. In other words, in these limits no finite surface element $\in \mathcal{S}$ would be missed by the random sampling.

The merits of this procedure are the following: 1) The random sampling on \mathcal{S} does not require further orbits to be calculated; 2) in contrast to acceptance-rejection methods common in Monte Carlo sampling, all points are accepted; and 3) in contrast to the brute force approach, no points are computed either interior or exterior to \mathcal{S} .

3.4. An example

To illustrate this calculation of confidence intervals, the model binary with elements given in Eq. (1) is observed in a campaign with parameters $f_{\text{orb}} = 0.6, N = 15, \sigma = 0''.05$. The initial 3D search for χ_{min}^2 is on a coarse 100^3 grid spanning the intervals $(1, 4), (0, 1), (0, 1)$ in $\log P, e, \tau$, respectively. Given the resulting initial estimate of χ_{min}^2 , the search boundaries are pruned in such a way that no point with $\chi^2 < \chi_{\text{min}}^2 + 25$ is excluded, and then a new 100^3 grid is computed. The resulting χ_{min}^2 and its location are then slightly improved by making small random displacements and accepting those that reduce χ^2 .

An investigator selects the confidence intervals of interest by specifying ν and α . Here we take $\alpha = 0.683$, corresponding to $\pm 1\sigma$ limits, and $\nu = 1$, thus computing confidence intervals for each element individually – i.e., independently of the other elements. With these choices, $\Delta_{\nu, \alpha} = 1$. For $\pm 2\sigma$ limits, $\alpha = 0.954$ and $\Delta_{\nu, \alpha} = 4$.

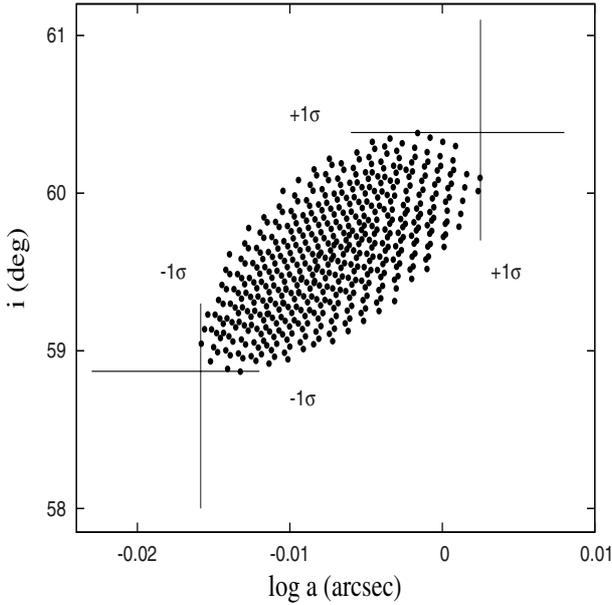


Fig. 1. Approximate confidence intervals. The ensemble of orbit vectors $\theta_{ijk} = (\phi_{ijk}, \hat{\psi}_{ijk})$ with $\hat{\chi}_{ijk}^2 < \chi_{\min}^2 + 1$ is projected on to the $(\log a, i)$ -plane. The $\pm 1\sigma$ bounds for each coordinate are indicated. The 506 grid points ϕ_{ijk} define the domain \mathcal{D} .

With $\Delta_{v,\alpha} = 1$, the refined grid has 506 points satisfying Eq. (5). These define \mathcal{D} and, as described in Sect. 3.2, approximate confidence intervals are obtained from projections of \mathcal{D} . The details are as follows: At each point $\phi_{ijk} \in \mathcal{D}$, the least squares values $\hat{A}, \hat{B}, \hat{F}, \hat{G}$ are available from the grid search. From these, the Campbell elements $\hat{a}, \hat{i}, \hat{\omega}, \hat{\Omega}$ are calculated as described in Sect. A4 of L14. Thus the Campbell elements θ_{ijk} are known at every point $\in \mathcal{D}$ and the projections of this ensemble give the $\pm 1\sigma$ intervals. Figure 1 illustrates this procedure. The $(\log a, i)$ -projection of the θ_{ijk} vectors $\in \mathcal{D}$ is plotted as are the resulting $\pm 1\sigma$ limits.

Now, for the same observed orbit $(\tilde{x}_n, \tilde{y}_n)$ and with the same refined grid, *accurate* $\pm 1\sigma$ intervals are computed according to the procedure of Sect. 3.3. At each of the 506 grid points $\in \mathcal{D}$, $\mathcal{N} = 5$ random points on \mathcal{S} are calculated as described in Sect. A.5. These points ψ_ℓ are such that $\chi^2(\psi_\ell) = \chi_{\min}^2 + 1$. For each ψ_ℓ , the corresponding Campbell elements are then derived as described in Sect. A.4 of L14. The final result is 2530 vectors $\theta_\ell \in \mathcal{S}$. The projections of \mathcal{S} give the desired $\pm 1\sigma$ limits. Figure 2 illustrates this step by again projecting the cloud of points on to the $(\log a, i)$ -plane.

Because Figs. 1 and 2 refer to the same orbit and are plotted to the same scale, we see immediately that, as anticipated in Sect. 3.2, the approximate $\pm 1\sigma$ intervals are enclosed by the accurate intervals. The ratios accurate:approximate are 1.4 for $\log a$ and 1.9 for i .

A point to note from Figs. 1 and 2 is that with finite samples, the derived confidence limits will always be underestimates. In the above example, increasing \mathcal{N} from 5 to 500 increases the confidence intervals for $\log a$ and i by additional factors of 1.045 and 1.055, respectively. Convergence experiments indicate that sufficient accuracy is achieved with $\mathcal{N} \gtrsim 200$.

The error bars derived with these procedures for a quantity Q are in general not symmetric about its minimum- χ^2 value. Accordingly, in testing these procedures, attention is focussed not on error bars but on confidence intervals (Q_L, Q_U) .

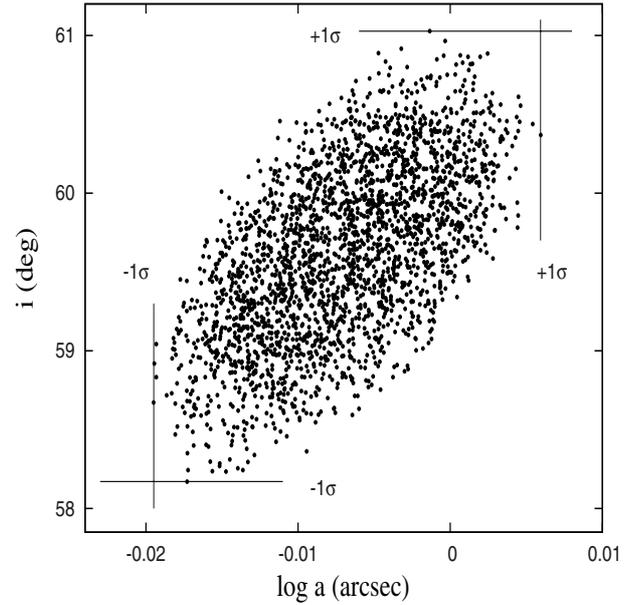


Fig. 2. Accurate confidence limits. The ensemble of orbit vectors $\theta_\ell \in \mathcal{S}$ – i.e., with $\chi_\ell^2 = \chi_{\min}^2 + 1$ – is projected on to the $(\log a, i)$ -plane. The $\pm 1\sigma$ bounds for each coordinate are indicated. Each of the 506 grid points $\phi_{ijk} \in \mathcal{D}$ gives rise to $\mathcal{N} = 5$ points on \mathcal{S} .

3.5. Coverage fractions

Confidence intervals calculated as described in Sect. 3.3 are claimed to be asymptotically rigorous. This implies that, with a fine enough grid and a sufficiently large \mathcal{N} , coverage fractions will be close to their frequentist values. To test this, two experiments similar to those in Sect. 5.5 of L14 are now reported. In each experiment, the example of Sect. 3.4 is repeated 10 000 times with independent realizations $(\tilde{x}_n, \tilde{y}_n)$ of the observed orbit. For each trial, confidence intervals are computed for various quantities Q , and an interval is counted as a success if $Q_L < Q_{\text{exact}} < Q_U$. The quantities Q are the Campbell elements plus the mass estimator \mathcal{M} defined in Eq. (17) of L14.

The approximate and accurate intervals of Sect. 3.2 and Sect. 3.3 – now with $\mathcal{N} = 500$ – are investigated in experiments I and II, respectively. In each experiment, both $\pm 1\sigma$ and $\pm 2\sigma$ intervals are tested.

The results are reported in Table 1. From experiment I, we see that the approximate confidence intervals are too small by factors of up to 1.9 for the $\pm 1\sigma$ intervals and up to 1.4 for the $\pm 2\sigma$ intervals. In contrast, the results for experiment II are close to ideal. Specifically, to within errors, the $\pm 2\sigma$ coverage fractions match the frequentist value. On the other hand, the $\pm 1\sigma$ coverage fractions fall short by the inconsequential factor of 1.016, indicating the need for a somewhat larger \mathcal{N} .

These results support the claim that the procedure developed in Sect. 3.3 is asymptotically rigorous. Moreover, the additional computational burden is negligible: experiment II required a mere 11% more computer time than experiment I.

4. Bayesian estimation

In this section, the Bayesian treatment of L14 is modified to eliminate its dependence on the profile likelihood – Eq. (3) in L14.

Table 1. Coverage fractions f for confidence intervals from 10^4 trials.

Q	Expt.	$\Delta_{v,\alpha} = 1$	$\Delta_{v,\alpha} = 4$
$\mathcal{E}(f)$		0.683 ± 0.005	0.954 ± 0.002
$\log P$	I	0.627 ± 0.005	0.948 ± 0.002
	II	0.674 ± 0.005	0.954 ± 0.002
e	I	0.646 ± 0.005	0.951 ± 0.002
	II	0.676 ± 0.005	0.956 ± 0.002
τ	I	0.639 ± 0.005	0.947 ± 0.002
	II	0.667 ± 0.005	0.953 ± 0.002
$\log a$	I	0.516 ± 0.005	0.859 ± 0.003
	II	0.670 ± 0.005	0.953 ± 0.002
i	I	0.382 ± 0.005	0.691 ± 0.005
	II	0.672 ± 0.005	0.952 ± 0.002
ω	I	0.589 ± 0.005	0.911 ± 0.003
	II	0.668 ± 0.005	0.954 ± 0.002
Ω	I	0.407 ± 0.005	0.736 ± 0.004
	II	0.675 ± 0.005	0.955 ± 0.002
$\log \mathcal{M}$	I	0.369 ± 0.005	0.684 ± 0.005
	II	0.672 ± 0.005	0.952 ± 0.002

4.1. Posterior density

The posterior probability density function (PDF) at (ϕ, ψ) given data D is

$$\Lambda(\phi, \psi|D) \propto \mathcal{L}(\phi, \psi|D) \pi(\phi, \psi) \quad (8)$$

where \mathcal{L} is the likelihood of (ϕ, ψ) given D , and $\pi(\phi, \psi)$ is a PDF that quantifies the investigator's prior beliefs or knowledge about (ϕ, ψ) . As in L14, we assume that π is the product of seven independent priors, one for each element. With the same choices as in L14, π can be omitted from Eq. (8) if the ϕ elements are now understood to be $(\log P, e, \tau)$.

Since coefficients independent of (ϕ, ψ) can be ignored

$$\mathcal{L}(\phi, \psi|D) \propto \exp\left(-\frac{1}{2}\hat{\chi}^2\right) \times \exp\left(-\frac{1}{2}\delta\chi^2\right) \quad (9)$$

where $\hat{\chi}^2(\phi) = \chi^2(\hat{\psi}|\phi)$ is the minimum value of χ^2 at fixed ϕ , and $\delta\chi^2(\delta\psi|\phi)$ is the positive increment in χ^2 for the displacement $\delta\psi = \psi - \hat{\psi}$. The second factor in Eq. (9) can be eliminated using Eq. (A.9), so that

$$\Lambda \propto \exp\left(-\frac{1}{2}\hat{\chi}^2\right) \eta(\phi|D) p(\psi|\phi, D). \quad (10)$$

If we now approximate $p(\psi|\phi, D)$ by a sum of δ functions as discussed in Sect. A.3, then

$$\Lambda \propto \exp\left(-\frac{1}{2}\hat{\chi}^2\right) \frac{\eta}{\mathcal{N}} \sum_{\ell} \delta(\psi - \psi_{\ell}) \quad (11)$$

where the ψ_{ℓ} are \mathcal{N} independent vectors that randomly sample the exact quadrivariate normal PDF.

The PDF Λ is for the Thiele-Innes elements. The corresponding PDF $\Gamma(\theta|D)$ for the Campbell elements θ is

$$\Gamma \propto \exp\left(-\frac{1}{2}\hat{\chi}^2\right) \frac{\eta}{\mathcal{N}} \sum_{\ell} \delta(\theta - \theta_{\ell}) \quad (12)$$

where $\theta_{\ell} = \vartheta(\psi_{\ell})$.

4.2. Credibility intervals

In terms of a 3D scan over ϕ -space, the PDF $\Gamma(\theta|D)$ is approximated by the ensemble of 7D vectors

$$\theta_m = (\phi_{ijk}, \vartheta_{\ell}) \quad (13)$$

with weights

$$\mu_m = \frac{\eta_{ijk}}{\mathcal{N}} \exp\left(-\frac{1}{2}\hat{\chi}_{ijk}^2\right) \quad (14)$$

Here m is an index that enumerates the random points ϑ_{ℓ} across all grid points (i, j, k) .

If $Q(\theta)$ is a quantity for which a credibility interval is required, the data from which this can be computed are the values $Q_m = Q(\theta_m)$ with weights μ_m . From this data, an estimate of the PDF of Q is

$$\Theta(Q) = \sum_m \mu_m \delta(Q - Q_m) / \sum_m \mu_m \quad (15)$$

with corresponding cumulative distribution function (CDF)

$$F(Q) = \sum_{Q_m < Q} \mu_m / \sum_m \mu_m. \quad (16)$$

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

$$F(Q_L) = 0.1587 \quad F(Q_U) = 0.8413 \quad (17)$$

so that the enclosed probability 0.6826.

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

4.3. Calculation procedure

The basic steps required to derive credibility intervals are as follows:

- 1) At every point $(\log P_i, e_j, \tau_k)$, the minimum- χ^2 Thiele-Innes elements $\hat{A}, \hat{B}, \hat{C}, \hat{G}$ are obtained with Eqs. (A.7) of L14, and the corresponding $\hat{\chi}_{ijk}^2$ computed.
- 2) The variances and covariances defining the exact quadrivariate normal PDF $p(\psi|\phi, D)$ at ϕ_{ijk} are computed with Eqs. (A.9), (A.10) of L14.
- 3) Random points ψ_{ℓ} sampling the exact PDF $p(\psi|\phi_{ijk}, D)$ are computed as described in Sect. A.4.
- 4) The Campbell elements ϑ_{ℓ} corresponding to ψ_{ℓ} are computed as described in Sect. A.4 of L14.
- 5) The vectors θ_m are then $(\phi_{ijk}, \vartheta_{\ell})$ with weights μ_m given by Eq. (14).
- 6) Lastly, credibility intervals are derived from Q_m with the approximate CDF given in Eq. (16).

4.4. An example

To illustrate this procedure, credibility intervals are computed for the orbit $(\tilde{x}_n, \tilde{y}_n)$ discussed in Sect. 3.4. A scatter diagram analogous to Figs. 1 and 2 is not readily constructed because the points θ_m are not of equal weight. Instead, the confidence intervals derived as in Fig. 2 (but now with $\mathcal{N} = 500$) are compared with the credibility intervals derived from Eqs. (17).

The $\Delta_{v,\alpha} = 1$ confidence interval for $\log a$ is $(-0.020, 0.006)$, whereas the equal-tail 68.3% credibility interval is $(-0.019, 0.008)$. The corresponding intervals for i are $(58^{\circ}1, 61^{\circ}1)$ and $(58^{\circ}2, 61^{\circ}2)$, respectively.

Table 2. Coverage fractions for credibility intervals from 10^4 trials.

Q	Expt.	1σ	2σ
$\mathcal{E}(f)$		0.683 ± 0.005	0.954 ± 0.002
$\log P$	III	0.664 ± 0.005	0.950 ± 0.002
e	III	0.675 ± 0.005	0.956 ± 0.002
τ	III	0.680 ± 0.005	0.949 ± 0.002
$\log a$	III	0.680 ± 0.005	0.951 ± 0.002
i	III	0.685 ± 0.005	0.954 ± 0.002
ω	III	0.682 ± 0.005	0.948 ± 0.002
Ω	III	0.677 ± 0.005	0.948 ± 0.002
$\log \mathcal{M}$	III	0.680 ± 0.005	0.953 ± 0.002

In these calculations, $\mathcal{N} = 50$ for $\chi_{ijk}^2 < \chi_{\min}^2 + 21.85$ and =1 otherwise. The domain defined by this inequality corresponds to $\Delta_{\nu,\alpha}$ with $\alpha = 0.9973$ and $\nu = 7$, thus ensuring an accurate treatment of the wings of the posterior PDF's to beyond $\pm 2\sigma$. Convergence experiments indicate that sufficient accuracy is achieved with $\mathcal{N} \gtrsim 20$.

In contrast to confidence intervals derived from scatter plots such as Fig. 2, the credibility intervals calculated from Eqs. (17) are not biased. Convergence to the asymptote is therefore faster and sufficient accuracy is achieved with a smaller \mathcal{N} .

4.5. Coverage fractions

For comparison with Sect. 3.5 above and with Sect. 5.5 of L14, coverage fractions for credibility intervals are computed for 10 000 independent realizations of the observed orbit, and an interval is again counted as a success if $Q_L < Q_{\text{exact}} < Q_U$. The results of this experiment (III) are given in Table 2.

Because these are credibility not frequentist intervals, there is no rigorous asymptotic expectation that the frequentist fractions should be recovered. Nevertheless, these ideal fractions are closely matched and so the credibility intervals calculated according to Sect. 4.2 can be described as *well-calibrated* (Drawid 1982).

When Table 2 is compared to Table 1 in L14, we see that the previous low coverage fractions for the ψ -elements $\log a$, i , ω , Ω and for the derived quantity $\log \mathcal{M}$ are now replaced by fractions close to their frequentist values. This confirms the conjecture in L14 that the shortfall was due to the profile likelihood.

As in Sect. 3.5, statistical rigour is achieved with only a modest increase in the computational burden. Experiment III required 22% more computer time than experiment I.

When Bayesian estimates depend on *informative* priors, the PDF $\pi(\theta)$ may have a significant gradient at θ_* , the elements of a particular binary. A coverage experiment restricted to θ_* will then (correctly) deviate from the frequentist expectation. In coverage tests for such cases, each independent orbit \tilde{x}_n, \tilde{y}_n should also be for a random θ drawn from $\pi(\theta)$.

5. Comparison of estimates

The relative performance of minimum- χ^2 and Bayesian estimation in experiments II and III is summarized in Table 3. The means $\langle \delta Q \rangle$ and standard deviations $s_{\delta Q}$ of the residuals $\delta Q = Q_{\text{est}} - Q_{\text{exact}}$ are tabulated, where Q_{est} is either the minimum- χ^2 value of Q or its posterior mean, and Q_{exact} is given in Eq. (1).

Table 3. Comparison of residuals δQ .

Q	$\langle \delta Q \rangle$	$s_{\delta Q}$	$\langle \delta Q \rangle$	$s_{\delta Q}$
	min $-\chi^2$		Bayes	
$\log P$	0.0021	0.025	0.0063	0.027
e	0.0017	0.026	0.0045	0.027
τ	-0.0013	0.024	-0.0044	0.025
$\log a$	0.0017	0.016	0.0041	0.016
i	0.0079	1.53	0.127	1.52
ω	-0.0070	3.22	-0.37	3.32
Ω	0.0001	2.25	0.133	2.31
$\log \mathcal{M}$	0.0010	0.034	-0.0002	0.035

Table 2 shows that in this test the two estimation methodologies yield closely similar results. This is to be expected for a non-informative prior $\pi(\theta)$ with negligible gradient at θ_* .

6. Conclusion

This paper has addressed a technical issue in the statistical analysis of orbits: how to achieve statistical rigour while taking advantage of the linearity of a subset of the orbital elements. The coverage experiments reported in Sects. 3.5 and 4.5 demonstrate that statistical rigour is achieved for both minimum- χ^2 and Bayesian estimation. Moreover, the reported timings show an inconsequential increase in the computational burden.

The key to this success is that at each grid point the distribution of probability in Thiele-Innes space is known. For minimum- χ^2 estimation, this allows Monte Carlo sampling to be targeted (Sect. A.5) precisely on the constant χ^2 surface defining the desired confidence level. For Bayesian estimation, the known PDF allows Monte Carlo sampling in Thiele-Innes space to be concentrated (Sect. A.4) on the high probability domain enclosing the least-squares points $(\hat{A}, \hat{B}, \hat{F}, \hat{G})$. Moreover, in each case, the random sampling does not require additional orbits to be computed.

The approach developed here is not specific to visual binaries or to the Thiele-Innes elements. In principle, an analogous procedure can be constructed for any partially linear estimation problem.

Acknowledgements. The issue of error underestimation in hybrid problems was raised by the referee of the previous paper (L14) and was the direct stimulus of this investigation. This same referee provided useful comments on this paper.

References

- Dawid, A. P. 1982, J. Am. Stat. Assoc., 77, 605
 Eastman, J., Gaudi, B., & Agol, E. 2013, PASP, 125, 83
 Hartkopf, W. I., McAlister, H. A., & Franz, O. G. 1989, AJ, 98, 1014
 James, F. 2006, Statistical Methods in Experimental Physics (Singapore: World Scientific Publishing Co.)
 Lucy, L. B. 2014, A&A, 563, 126 (L14)
 Muller, M. E. 1959, Comm. Assoc. Comp. Mach., 2, 19
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical Recipes, 2nd edn. (Cambridge: Cambridge Univ. Press)
 Schaefer, G. H., Simon, M., Beck, T. L., Nelan, E., & Prato, L. 2006, AJ, 132, 2618

Appendix A: Statistics in Thiele-Innes space

In Appendix A of L14, formulae are derived for $\hat{\psi} = (\hat{A}, \hat{B}, \hat{F}, \hat{G})$, the least squares Thiele-Innes constants at given $\phi = (\log P, e, \tau)$. The resulting value of χ^2 then determines the profile likelihood \mathcal{L}^\dagger used in the approximation of posterior means – Eq. (7) of L14. Now we wish to sample points displaced from $\hat{\psi}$. Let such a displacement be $\delta\psi = (a, b, f, g)$.

A.1. Probability density function $p(\psi|\phi, D)$

In Sect. A.4 of L14, the PDF at $\delta\psi$ is shown to be the product of two *independent* PDF's, each a bivariate normal distribution, one for (a, f) , the other for (b, g) . Formulae for the variances $\sigma_a^2, \sigma_b^2, \sigma_f^2, \sigma_g^2$ and the covariances $\text{cov}(a, f), \text{cov}(b, g)$ that define these PDF's are given in Eqs. (A.9) and (A.10) of L14.

The PDF for (a, f) is

$$p(a, f) = \frac{1}{2\pi\sigma_a\sigma_f\sqrt{1-\rho_{af}^2}} \exp\left(-\frac{1}{2}\xi^2\right) \quad (\text{A.1})$$

where $\rho_{af} = \text{cov}(a, f)/(\sigma_a\sigma_f)$ and

$$\xi^2 = \frac{1}{1-\rho_{af}^2} \left[\left(\frac{a}{\sigma_a}\right)^2 + \left(\frac{f}{\sigma_f}\right)^2 - 2\rho_{af} \left(\frac{a}{\sigma_a}\right)\left(\frac{f}{\sigma_f}\right) \right] \quad (\text{A.2})$$

The point $(0, 0)$ corresponds to the minimum- χ_x^2 solution (\hat{A}, \hat{F}) , where χ_x^2 is the x -coordinate contribution to χ^2 – see Eq. (2). The displacement (a, f) therefore results in a positive increment $\delta\chi_x^2$ given by

$$\sigma^2\delta\chi_x^2 = \sum_n w_n [(x_n - \bar{x}_n)^2 - (\hat{x}_n - \bar{x}_n)^2]. \quad (\text{A.3})$$

Now, for displacement (a, f) , the predicted

$$x_n = \hat{x}_n + a X_n + f Y_n \quad (\text{A.4})$$

where X, Y are given by Eqs. (A.3) in L14. Substitution in Eq. (A.3) then gives

$$\sigma^2\delta\chi_x^2 = \sum_n w_n (a^2 X_n^2 + f^2 Y_n^2 + 2af X_n Y_n) \quad (\text{A.5})$$

where terms linear in a and f vanish because the minimum is at $(0, 0)$. The summations in Eq. (A.5) can be eliminated in favour of σ_a^2, σ_f^2 and $\text{cov}(a, f)$ with the formulae given in Eqs. (A.6), (A.9) and (A.10) of L14. After lengthy algebra, we find that

$$\delta\chi_x^2 = \frac{1}{1-\rho_{af}^2} \left[\left(\frac{a}{\sigma_a}\right)^2 + \left(\frac{f}{\sigma_f}\right)^2 - 2\rho_{af} \left(\frac{a}{\sigma_a}\right)\left(\frac{f}{\sigma_f}\right) \right] \quad (\text{A.6})$$

and so

$$p(a, f) \propto \exp\left(-\frac{1}{2}\delta\chi_x^2\right). \quad (\text{A.7})$$

Exactly the same analysis applies to the independent pair (b, g) , so that

$$p(b, g) \propto \exp\left(-\frac{1}{2}\delta\chi_y^2\right). \quad (\text{A.8})$$

Combining these formulae, we find that the PDF at $\psi = \hat{\psi} + \delta\psi$ is

$$p(\psi|\phi, D) = \frac{1}{4\pi^2\sigma^4\eta} \exp\left(-\frac{1}{2}\delta\chi^2\right) \quad (\text{A.9})$$

where $\delta\chi^2 = \delta\chi_x^2 + \delta\chi_y^2$ and $\eta(\phi|D)$ is given by

$$\sigma^4\eta = \sigma_a\sigma_f\sqrt{1-\rho_{af}^2} \times \sigma_b\sigma_g\sqrt{1-\rho_{bg}^2} \quad (\text{A.10})$$

A.2. Modified Thiele-Innes constants

The familiar device of “completing the square” applied to Eq. (A.2) suggests new variables \mathcal{A}, \mathcal{F} defined by

$$\frac{a}{\sigma_a} = \mathcal{A}, \quad \frac{f}{\sigma_f} = \rho_{af}\mathcal{A} + \sqrt{1-\rho_{af}^2}\mathcal{F}. \quad (\text{A.11})$$

Substitution in Eq. (A.6) then gives

$$\delta\chi_x^2 = \mathcal{A}^2 + \mathcal{F}^2. \quad (\text{A.12})$$

The Jacobian of this transformation is

$$\frac{\partial(a, f)}{\partial(\mathcal{A}, \mathcal{F})} = \sigma_a\sigma_f\sqrt{1-\rho_{af}^2} \quad (\text{A.13})$$

so that, by conservation of probability, the PDF of $(\mathcal{A}, \mathcal{F})$ is

$$\Psi(\mathcal{A}, \mathcal{F}) = \frac{1}{2\pi} \exp\left[-\frac{1}{2}(\mathcal{A}^2 + \mathcal{F}^2)\right]. \quad (\text{A.14})$$

Now, exactly the same analysis applies to the independent pair (b, g) . Thus, if we define new variables \mathcal{B}, \mathcal{G} by the equations

$$\frac{b}{\sigma_b} = \mathcal{B}, \quad \frac{g}{\sigma_g} = \rho_{bg}\mathcal{B} + \sqrt{1-\rho_{bg}^2}\mathcal{G} \quad (\text{A.15})$$

then the PDF of $(\mathcal{B}, \mathcal{G})$ is $\Psi(\mathcal{B}, \mathcal{G})$, and

$$\delta\chi_y^2 = \mathcal{B}^2 + \mathcal{G}^2. \quad (\text{A.16})$$

If ζ denotes the vector $(\mathcal{A}, \mathcal{B}, \mathcal{F}, \mathcal{G})$, then $\Pi(\zeta)$, the PDF in ζ -space, is $\Psi(\mathcal{A}, \mathcal{F}) \times \Psi(\mathcal{B}, \mathcal{G})$ – i.e.,

$$\Pi(\zeta) = \frac{1}{4\pi^2} \exp\left[-\frac{1}{2}(\mathcal{A}^2 + \mathcal{B}^2 + \mathcal{F}^2 + \mathcal{G}^2)\right]. \quad (\text{A.17})$$

Accordingly, the distribution of probability in ζ -space is simply the product of four independent normal distributions, each with zero mean and unit variance.

The increment in χ^2 is given by Eqs. (A.12) and (A.16) as

$$\delta\chi^2(\zeta) = \mathcal{A}^2 + \mathcal{B}^2 + \mathcal{F}^2 + \mathcal{G}^2. \quad (\text{A.18})$$

A.3. Approximate PDFs

A distribution of probability can be represented by a sum of δ functions in such a way that the probability attached to any *finite* element of space is approximated with arbitrary accuracy. Thus, the PDF giving the distribution of probability in ζ -space is given approximately by

$$p(\zeta) = \mathcal{N}^{-1} \sum_{\ell} \delta(\zeta - \zeta_{\ell}) \quad (\text{A.19})$$

where each ζ_{ℓ} is an independent random vector sampling the PDF $\Pi(\zeta)$ given by Eq. (A.17). The integral of $p(\zeta)$ over a finite element in ζ -space converges to the exact value as $\mathcal{N} \rightarrow \infty$.

Equations (A.11) and (A.15) transform the point ζ_{ℓ} into the displacement $\delta\psi_{\ell} = \psi_{\ell} - \hat{\psi}$. The corresponding approximate PDF in ψ -space is therefore

$$p(\psi|\phi, D) = \mathcal{N}^{-1} \sum_{\ell} \delta(\psi - \psi_{\ell}). \quad (\text{A.20})$$

Note that the Jacobian of this transformation from ζ - to ψ -space is implicit in the changes in number densities of the delta functions in the respective spaces.

Similarly, the PDF for the Campbell elements ϑ corresponding to the Thiele-Innes elements ψ is

$$p(\vartheta|\phi, D) = \mathcal{N}^{-1} \sum_{\ell} \delta(\vartheta - \vartheta_{\ell}) \quad (\text{A.21})$$

where $\vartheta_{\ell} = \vartheta(\psi_{\ell})$ is derived as described in Sect. A.4 of L14.

A.4. Random sampling in ψ -space

According to Eq. (A.17), a random point in ζ -space is (z_1, z_2, z_3, z_4) , where the z_i are independent random Gaussian variates drawn from $\mathcal{N}(0, 1)$. This point corresponds to the displacement (a, b, f, g) given by Eqs. (A.11) and (A.15) and therefore to the point $(\hat{A} + a, \hat{B} + b, \hat{F} + f, \hat{G} + g)$ in ψ -space. Thus, a point randomly selected from the exact PDF $p(\psi|\phi, D)$ can be derived from four independent Gaussian variates, and the resulting increment in χ^2 is given by Eq. (A.18) as

$$\delta\chi^2(\psi) = z_1^2 + z_2^2 + z_3^2 + z_4^2 \quad (\text{A.22})$$

A.5. Random sampling at fixed $\delta\chi^2$

A random point in ψ -space subject to a constraint on $\delta\chi^2$ can be found by first selecting a random point on the 4D sphere in ζ -space defined by Eq. (A.18). This is achieved as follows: If z_i again denotes a Gaussian variate from $\mathcal{N}(0, 1)$, then a random point on this hypersphere is $(z_1, z_2, z_3, z_4)/Z$, where

$$Z^2 = (z_1^2 + z_2^2 + z_3^2 + z_4^2) / \delta\chi^2 \quad (\text{A.23})$$

(Muller 1979). The corresponding point (A, B, F, G) in ψ -space is then derived from Eqs. (A.11) and (A.15).

The random sampling procedures of Sects. A.4 and A.5 *predict* χ^2 without the need to compute an orbit. This is achieved by exploiting the linearity of the ψ -elements and is the basis of the computational efficiency of the techniques of Sects. 3.3 and 4.1. However, during code development, this prediction should be tested by actually computing the orbit and independently evaluating χ^2 from Eq. (2).