Sampling errors in nested sampling parameter estimation

Edward Higson^{*,¶,||}, Will Handley^{†,¶,||}, Mike Hobson^{‡,¶} and Anthony Lasenby^{§,¶,||}

Abstract. Sampling errors in nested sampling parameter estimation differ from those in Bayesian evidence calculation, but have been little studied in the literature. This paper provides the first explanation of the two main sources of sampling errors in nested sampling parameter estimation, and presents a new diagrammatic representation for the process. We find no current method can accurately measure the parameter estimation errors of a single nested sampling run, and propose a method for doing so using a new algorithm for dividing nested sampling runs. We empirically verify our conclusions and the accuracy of our new method.

Keywords: nested sampling, parameter estimation.

1 Introduction

Nested sampling (Skilling, 2006) is a Monte Carlo method for Bayesian analysis which simultaneously calculates both Bayesian evidences and posterior samples. The early development of the algorithm was focused on evidence calculation, which is computationally expensive using variants of standard MCMC sampling based on the Metropolis-Hastings algorithm (MacKay, 2003).

Contemporary implementations such as MULTINEST (Feroz and Hobson, 2008; Feroz et al., 2009, 2013) and POLYCHORD (Handley et al., 2015a,b) are now also extensively used for parameter estimation from posterior samples (see for example Planck Collaboration, 2016). Nested sampling compares favourably to MCMC-based parameter estimation for degenerate, multi-modal likelihoods as it has no "thermal" transition probability and exponentially compresses the prior distribution to the posterior. However, despite its increasing popularity, the sampling errors in nested sampling parameter estimation are poorly understood.

Correctly quantifying uncertainty is vital for identifying spurious results — in particular we find sampling errors often significantly affect estimates of credible intervals on parameters. Conversely, finding such errors are very small may imply an unnecessarily large amount of computational resource is being used for the calculation. This paper has two goals: to provide an explanation of the sources of these errors and an

© 0000 International Society for Bayesian Analysis

^{*}e.higson@mrao.cam.ac.uk

[†]wh260@mrao.cam.ac.uk

[‡]mph@mrao.cam.ac.uk

[§]a.n.lasenby@mrao.cam.ac.uk

[¶]Astrophysics Group, Battcock Centre, Cavendish Laboratory, JJ Thomson Avenue, Cambridge CB3 0HE, UK

Kavli Institute for Cosmology, Madingley Road, Cambridge, CB3 0HA, UK

empirical technique for estimating them. One obvious method is to repeat the analysis some n_{repeats} times, although this increases the computational cost by a corresponding factor. Interestingly, we find no current method can accurately estimate these errors on parameter estimates from a single analysis, and so we present a new method for doing this. Our approach uses a new algorithm for dividing a single nested sampling run into multiple valid nested sampling runs; these can then be recombined in different combinations using resampling techniques such as the bootstrap. We test our results and new method empirically.

The paper begins with background on sampling errors in parameter estimation from posterior samples, then describes the nested sampling algorithm and how it is currently used for parameter estimation in Section 2. We explain the two main sources of sampling errors in nested sampling parameter estimation in Section 3, and present a new diagrammatic representation of the process (illustrated in Figures 3(a) to 3(e)). Section 4 describes our new method for measuring sampling errors from a single nested sampling run, using our new algorithm for division of such runs.

We empirically test our method's accuracy in Section 5 with the help of analytical cases in the manner described by Keeton (2011). Here one can obtain uncorrelated samples from the prior space within some likelihood contour using standard techniques, and we term the resulting procedure *perfect nested sampling*. Our approach accurately quantifies uncertainties on parameter estimates from the stochasticity of the nested sampling algorithm, but software used for practical problems may produce additional errors from correlated samples within likelihood contours that are specific to a given implementation. We discuss implementation-specific errors in Section 6, including testing sampling error estimates from our method for POLYCHORD calculations. Our method gives superior performance to the current approach and can be easily included in nested sampling software; we are currently working on incorporating it into future versions of POLYCHORD.

Background: sampling errors in parameter estimation

Sampling can be used to represent a posterior distribution $\mathcal{P}(\theta)$ via a set of weighted samples

$$\mathcal{S} = \{(\theta_s, p_s), s = 1, \dots, n_{\text{samp}}\},\tag{1}$$

where each θ_s is drawn from the posterior distribution with probability proportional to $p_s \times \mathcal{P}(\theta_s)$, and $\sum_{s \in S} p_s = 1$. Likelihoods \mathcal{L} are often computationally expensive functions, so the goal of parameter estimation is to sample the posterior distribution $\mathcal{P}(\theta)$ numerically with a limited number of likelihood calls.

Samples S may be used to compute numerical results. For example, the posterior expectation of a function of the parameters $f(\theta)$ can be estimated as

$$\mathbf{E}[f(\theta)] = \int f(\theta) \mathcal{P}(\theta) \, d\theta \approx \sum_{s \in \mathcal{S}} p_s f(\theta_s).$$
(2)

In this case the sampling error is the difference between $\sum_{s \in S} p_s f(\theta_s)$ and the exact value of $E[f(\theta)]$. Often the posterior distributions of parameters θ are of interest, and

are estimated numerically from the samples by dividing the parameter space into cells or via kernel density estimation.

There have been many works on approximating MCMC sampling errors, including investigation of quantiles and the amount of computation required to reach some level of accuracy — see for example Doss et al. (2015); Flegal et al. (2008); Liu et al. (2016). In particular Sequential Monte Carlo samplers (Del Moral et al., 2006) have similarities with nested sampling, and their sampling errors are better understood. For some related methods such as the Tootsie Pop algorithm (Huber and Schott, 2014) and accelerated simulated annealing (Bezáková et al., 2008) the error distribution is known exactly, although these techniques are less widely used. This paper introduces empirically tested techniques for quantifying sampling errors from the nested sampling algorithm.

2 The nested sampling algorithm

Nested sampling (Skilling, 2006) is a numerical method computing Bayesian evidences

$$\mathcal{Z} = \int \mathcal{L}(\theta) \pi(\theta) \,\mathrm{d}\theta \tag{3}$$

and samples from the posterior distribution

$$\mathcal{P}(\theta) = \frac{\mathcal{L}(\theta)\pi(\theta)}{\mathcal{Z}} \tag{4}$$

given some likelihood $\mathcal{L}(\theta)$ and prior $\pi(\theta)$.

Initially *n* points, termed *live points*, are sampled randomly from the prior. At each iteration *i*, the live point with the lowest likelihood \mathcal{L}_i is removed and replaced by a new live point sampled from the prior subject to the constraint that it has a likelihood higher than \mathcal{L}_i . Iterating until some termination condition is met generates a list of discarded samples known as *dead points*, which are used to estimate the evidence and make posterior inferences¹. We refer to the completed nested sampling process as a *run*.

To compute the evidence, the many-dimensional integral (3) is reduced to a onedimensional integral in terms of the fractional prior volume within an iso-likelihood contour. We define the fraction of the prior θ with likelihood $\mathcal{L}(\theta)$ greater than some value \mathcal{L}^* as $X(\mathcal{L}^*)$, where

$$X(\mathcal{L}^*) \equiv \int_{\mathcal{L}(\theta) > \mathcal{L}^*} \pi(\theta) \,\mathrm{d}\theta,\tag{5}$$

and $X \in [0,1]$. Provided the inverse $\mathcal{L}(X) \equiv X^{-1}(\mathcal{L})$ exists², the evidence (3) can be expressed as

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X) \, \mathrm{d}X. \tag{6}$$

¹The remaining live points at termination can also be used if required, but termination conditions can be chosen such that this makes a negligible difference to calculation results.

²A sufficient condition for $\mathcal{L}(X) \equiv X^{-1}(\mathcal{L})$ to exist is for \mathcal{L} to be continuous and π to have a connected support. See Chopin and Robert (2010) and Feroz et al. (2013, Appendix C) for a more detailed measure-theoretic discussion.



Figure 1: A schematic representation of nested sampling with a constant number of live points n. The curve $\mathcal{L}(X)X$ shows the relative posterior mass, the bulk of which is contained in some small fraction $\exp(-H)$ of the prior and is only visible on a log scale in X. The algorithm iterates inwards in X exponentially with stochastic shrinkage ratios distributed according to (7).

Given a set of dead points with likelihoods \mathcal{L}_i , the corresponding prior volumes X_i are unknown but are modelled statistically as $X_i = t_i X_{i-1}$, where $X_0 = 1$ and each shrinkage ratio t_i is independently distributed as the largest of n random variables from the interval [0, 1] (Skilling, 2006). Hence:

$$P(t_i) = nt_i^{n-1}, \qquad \text{E}[\log t_i] = -\frac{1}{n}, \qquad \text{Var}[\log t_i] = \frac{1}{n^2},$$
 (7)

and the algorithm samples within an exponentially shrinking part of the prior. This exponential shrinkage is shown schematically in Figure 1.

2.1 Evidence estimation

Nested sampling therefore allows one to approximate the evidence (6) via a quadrature sum over the dead points

$$\mathcal{Z}(\mathbf{t}) \approx \sum_{i \in \text{dead}} \mathcal{L}_i w_i(\mathbf{t}),$$
 (8)

where $\mathbf{t} = \{t_1, t_2, \ldots, t_{n_{\text{dead}}}\}\$ are the unknown set of shrinkage ratios for the $n_{\text{dead}}\$ iterations of the nested sampling process, and each t_i is an independent random variable drawn from distribution (7). The shrinkage ratios define the prior volumes via $X_i(\mathbf{t}) = \prod_{k=0}^{i} t_k$, and the w_i are appropriately chosen quadrature weights roughly corresponding to the volume of the "prior shell" to which a given dead point belongs. For example, using the trapezium rule: $w_i(\mathbf{t}) = \frac{1}{2}(X_{i-1}(\mathbf{t}) - X_{i+1}(\mathbf{t}))$.

Given that the shrinkage ratios \mathbf{t} are a priori unknown, we may quantify our knowledge of \mathcal{Z} by simulating sets of \mathbf{t} according to (7), and working with the distribution

of the resulting set of evidences $\{\mathcal{Z}\}_{\mathbf{t}}$ from (8) (Skilling, 2006). Typically one then computes and reports a mean value and error for log \mathcal{Z} from this distribution.

Several alternative methods for calculating evidence inferences are reported in the literature. Skilling (2006) also proposes an error calculation based on relative entropy, which demonstrates that the uncertainty of $\log Z$ is dominated by the Poisson variability in the number of steps required to reach the bulk of the posterior mass. Keeton (2011) uses distribution moments and running totals which are updated with each nested sampling step. This method has been extended by Handley et al. (2015b) to allow the splitting of multi-modal likelihoods into different clusters and the treatment of variable numbers of live points. For a more detailed discussion of the convergence properties of nested sampling evidences, see Chopin and Robert (2010).

Thus, the dominant sampling error in the evidence estimate (8) from perfect nested sampling is from statistical variation in the unknown volumes of the prior "shells" $w_i(\mathbf{t})$ that each point represents. The error from approximating the integral for \mathcal{Z} with a sum can be safely neglected unless n is very small³ (Skilling, 2006). There is also some error from terminating the algorithm and truncating the sum, but this is can be made negligible with appropriate termination conditions.

2.2 Parameter estimation

One may also perform posterior inference from nested sampling by using the dead points to construct a set of posterior samples with weights proportional to their share of the posterior mass (Skilling, 2006):

$$p_i(\mathbf{t}) = \frac{w_i(\mathbf{t})\mathcal{L}_i}{\sum_i w_i(\mathbf{t})\mathcal{L}_i} = \frac{w_i(\mathbf{t})\mathcal{L}_i}{\mathcal{Z}(\mathbf{t})}.$$
(9)

As before, **t** is the set of prior shrinkage ratios and in the trapezium rule case $w_i(\mathbf{t}) = \frac{1}{2}(X_{i-1}(\mathbf{t}) - X_{i+1}(\mathbf{t}))$.

The weights defined by (9) present a departure from traditional sampling approaches in that the $w_i(\mathbf{t})$ are random variables, with their stochasticity determined by (7). When computing expectations (2) there is now an additional error associated with our lack of knowledge of the precise values $p_i(\mathbf{t})$. Nested sampling software packages such as MULTINEST and POLYCHORD produce posterior files containing only the expected values

$$\mathbf{E}[p_i(\mathbf{t})] = \frac{\mathrm{e}^{-i/n} \mathcal{L}_i}{\sum_j \mathrm{e}^{-j/n} \mathcal{L}_j}.$$
(10)

To account for the stochasticity in the weights p_i , Skilling (2006) suggests simulating the prior volume shrinkage ratios **t** in the same manner as for evidence estimation, and using these simulations to calculate a set of values for estimators such as (2). The sampling error should then be estimated from the variation within this sample; we term

³The trapezium rule error is $\mathcal{O}(1/n^2)$, and if required other methods such as Simpson integration could be used.

this the *simulated weights method*. We believe this procedure is the only estimate of sampling errors in parameter estimation from a single nested sampling run proposed in the literature. However it is in general an *underestimate*, as can be seen in the numerical tests in Section 5. Section B of the supplementary material discusses this underestimation of errors in detail.

We now describe why the simulated weights method does not capture all sources of sampling errors, and in Section 4 we propose a new method for correctly computing these errors.

3 Sources of sampling errors in nested sampling parameter estimation

In order to understand why the simulated weights method underestimates sampling errors, we require a result from Chopin and Robert (2010). They show that the expectation integral (2) may be re-phrased in terms of the prior volume X via:

$$\mathbf{E}[f(\theta)] = \int f(\theta)\mathcal{P}(\theta) \, d\theta = \int f(\theta)\frac{\mathcal{L}(\theta)\pi(\theta)}{\mathcal{Z}} \, \mathrm{d}\theta = \frac{1}{\mathcal{Z}}\int \tilde{f}(X)\mathcal{L}(X) \, \mathrm{d}X, \quad (11)$$

where $\tilde{f}(X)$ is the prior expectation of $f(\theta)$ on some iso-likelihood contour $\mathcal{L}(\theta) = \mathcal{L}(X)$,

$$\tilde{f}(X) \equiv \mathbf{E}^{\pi}[f(\theta)|\mathcal{L}(\theta) = \mathcal{L}(X)].$$
(12)

The simulated weights approach amounts to discretising the integral (11) as

$$\frac{1}{\mathcal{Z}} \int \tilde{f}(X) \mathcal{L}(X) \, \mathrm{d}X \approx \frac{1}{\mathcal{Z}} \sum_{i} \tilde{f}(X_i) \, \mathcal{L}_i \, \frac{1}{2} (X_{i-1} - X_{i+1}), \tag{13}$$

and, most importantly, further requiring that we may use $f(\theta_i)$ as a proxy for $\tilde{f}(X_i)$ at each point X_i , In some special cases $f(\theta_i) = \tilde{f}(X_i)$ for all θ and this approach is valid, for example when $f(\theta_i) = \tilde{f}(X_i) \propto -\log \mathcal{L}_i$ (entropy computation), but in general it is not. This can cause significant inaccuracies as iso-likelihood contours often span wide ranges of different parameter values, as illustrated in Figure 2.

To summarise, the dominant sampling errors in estimating some parameter or function of parameters from perfect nested sampling typically come from two sources:

- (i) approximating the unknown prior volumes $w_i(\mathbf{t})$ with their expectation $\mathbf{E}[w_i(\mathbf{t})]$ using (7);
- (ii) approximating the mean value of a function of parameters over an entire isolikelihood contour $\tilde{f}(X_i)$ with its value at a single point $f(\theta_i)$.

Errors from (i) are also present in evidence calculation; in the parameter estimation case they are typically smaller as results depend only on the *relative* weights of the



Figure 2: Nested sampling dead points and iso-likelihood contours for two-dimensional а multi-modal likelihood $\mathcal{L}(\theta)$; darker shading shows higher likelihoods. Isolikelihood contours can pass through a wide range of different parameter values.

samples. In contrast (ii) is only present in parameter estimation, where it is typically a significant or dominant source of sampling errors. The relative contributions of (i) and (ii) are empirically tested in Section A of the supplementary material, where they are calculated for analytical cases by using exact values for weights $w_i(\mathbf{t})$ and by replacing $f(\theta_i)$ with $\tilde{f}(X_i)$. The simulated weights method underestimates sampling errors in parameter estimation as it ignores errors from (ii).

We now introduce a new diagrammatic representation of nested sampling parameter estimation to illustrate the two different sources of sampling errors.

3.1 Diagrammatic representation

Nested sampling transforms evidence calculations of any dimension into a 1-dimensional problem⁴ in $\mathcal{L}(X)$ which can be entirely represented on a diagram like Figure 1. An analogous diagram for parameter estimation must also illustrate sampling a single point $f(\theta_i)$ on each iso-likelihood contour $\mathcal{L}(\theta) = \mathcal{L}(X_i)$ from the distribution $P(f(\theta)|X_i)$.

We propose a generalisation of Figure 1 for visualising parameter estimation problems, and present it in Figures 3(a) to 3(e). The top panel in each figure is similar to Figure 1 and shows the relative posterior mass $\mathcal{L}(X)X$ at each value of log X. The lower central panel shows the probability distribution $P(f(\theta)|X)$ and its mean $\tilde{f}(X)$. The posterior distribution is shown on the left — this is equal to the distributions $P(f(\theta)|X)$ (the lower central panel) marginalised over X in proportion to the posterior weight at each X (the top panel).

For these example plots we use d-dimensional spherical unit Gaussian likelihoods

$$\mathcal{L}(\theta) = (2\pi)^{-d/2} e^{-|\theta|^2/2}$$
(14)

⁴For practical nested sampling problems implementation-specific errors can differ for two likelihoods with the same $\mathcal{L}(X)$. For example if one likelihood has a much higher dimension and a much larger number of modes than the other it may have larger errors from the implementation software.







(b) $f(\theta) = \theta_{\hat{1}}$ with a 5-dimensional Cauchy likelihood (15) and a Gaussian prior (16).



(c) $f(\theta) = \theta_1$ with a 3-dimensional Gaussian likelihood (14) and a Gaussian prior (16).





(e) $f(\theta) = |\theta|$ (i.e. the radial distance from the likelihood's maximum) with a 5-dimensional Gaussian likelihood (14) and a Gaussian prior (16). In this case $f(\theta_i) = \tilde{f}(X_i)$ for all θ and sampling errors are only from uncertainty in prior volume shrinkages and the trapezium rule approximation.

Figure 3: Nested sampling parameter estimation diagrams: in each case the top panel shows the relative posterior mass at each value of $\log X$ ($\propto \mathcal{L}(X)X$). The lower central panel shows the distribution $P(f(\theta)|X)$ of values $f(\theta)$ on each iso-likelihood contour $\mathcal{L}(\theta) = \mathcal{L}(X)$; the dashed line shows the expectation of this distribution which we defined in (12) as $\tilde{f}(X)$. The left panel shows the posterior distribution of $f(\theta)$, with the dotted line showing its posterior expectation. The colour scale shows the fraction of the cumulative probability distribution lying between some region and the median. and d-dimensional spherical unit Cauchy likelihoods

$$\mathcal{L}(\theta) = \frac{\Gamma(\frac{1+d}{2})}{\pi^{(d+1)/2}} \left(1 + |\theta|^2\right)^{-(\frac{d+1}{2})},\tag{15}$$

with d-dimensional co-centred spherical Gaussian priors

$$\pi(\theta) = (2\pi\sigma_{\pi}^2)^{-d/2} e^{-|\theta|^2/2\sigma_{\pi}^2}, \qquad \sigma_{\pi} = 10.$$
(16)

We denote the first component of the θ vector as $\theta_{\hat{1}}$, although by symmetry the results will be the same for any component. $\overline{\theta_{\hat{1}}}$ and $\overline{\theta_{\hat{1}}}^2$ are the first and second moments of the posterior distribution of $\theta_{\hat{1}}$.

The form of the distribution $P(f(\theta)|X)$ as X varies depends on the likelihood only through the shape of the iso-likelihood contours $\mathcal{L}(\theta) = \mathcal{L}(X)$. Therefore the lower central panel of the diagrams for some $f(\theta)$ is the same for any likelihoods with the same contours — this can be seen in Figures 3(a) and 3(b), where the differences in the posterior (left panel) are due only to the different posterior weights in log X (top panel).

These diagrams can be constructed for any nested sampling calculation using the posterior samples and kernel density estimation; this could provide insight into the nature of the calculation and the relative contributions from the different sources of sampling errors.

3.2 Transforming a parameter estimation problem into 2 dimensions

As illustrated by our diagrams, nested sampling parameter estimation is fundamentally a 2-dimensional problem in $\mathcal{L}(X)$ and $P(f(\theta)|X)$. In fact a parameter estimation calculation for some $f(\theta)$ given $\mathcal{L}(\theta)$ is equivalent to a 2-dimensional problem for $f^*(\theta^*)$ given $\mathcal{L}^*(\theta^*)$ when

$$\mathcal{L}^*(\theta^*) = \mathcal{L}(X), \tag{17}$$

$$P(f^*(\theta^*)|X) = P(f(\theta)|X), \tag{18}$$

for all X. Any transformation satisfying (17) and (18) will leave our proposed diagram for the calculation unchanged. Parameter estimation can also be represented as a 1dimensional problem in $\mathcal{L}^*(\theta^*) = \mathcal{L}(X)$ combined with a univariate stochastic process for each dead point *i* with the distribution $P(f(\theta)|X_i)$.

One way to express a general nested sampling calculation in 2 dimensions is to map it onto the unit square $\theta^* = (X, Y)$ with uniform priors $X, Y \in [0, 1]$ and a likelihood $\mathcal{L}^*(\theta^*) = \mathcal{L}(X)$ which is independent of Y and satisfies (17). In this case X is as before the remaining fractional prior volume and Y parameterises each iso-likelihood contour. Using inverse transform sampling, for a general $f(\theta)$ a corresponding $f^*(\theta^*)$ satisfying (18) is

$$f^*(\theta^*) = f^*(X, Y) = F^{-1}(Y|X), \tag{19}$$

where $F^{-1}(Y|X)$ is the inverse of the cumulative distribution

$$F(Y|X) = \int_{-\infty}^{Y} P(f(\theta) = h|X) \,\mathrm{d}h.$$
⁽²⁰⁾

As an example let us consider *d*-dimensional spherically symmetric likelihoods such as (14) or (15) with co-centred spherically symmetric priors such as (16). Then $X(\theta)$ is a function only of the radial distance from the centre $|\theta|$, and the iso-likelihood contours $\mathcal{L}(\theta) = \mathcal{L}(X)$ are hyperspherical shells of some radius $|\theta_X|$. The probability distribution of a single parameter θ_1 (a single component of θ) on such an iso-likelihood contour is then

$$P(\theta_{\hat{1}}|X) = \begin{cases} \frac{\Gamma(\frac{d}{2})}{|\theta_X| \, \Gamma(\frac{1}{2}) \, \Gamma(\frac{d-1}{2})} \left(1 - \frac{\theta_{\hat{1}}^2}{|\theta_X|^2}\right)^{\frac{d-3}{2}} & \text{if } - |\theta_X| < \theta_{\hat{1}} < |\theta_X|, \\ 0 & \text{otherwise.} \end{cases}$$
(21)

 $\theta_{\hat{1}}$ can be sampled directly or used to calculate the inverse cumulative distribution which together with knowledge of the function $\mathcal{L}(X)$ allows the parameter estimation of a *d*-dimensional Gaussian to be transformed into a 2-dimensional problem on the unit square.

Samples from (21) can be generated efficiently using the symmetry around $\theta_{\hat{1}} = 0$ and the change of variables $\Theta = \theta_{\hat{1}}^2 / |\theta_X|^2$ to give a Beta distribution

$$P(\Theta|X) = \begin{cases} \frac{\Gamma(\frac{d}{2})}{\Gamma(\frac{1}{2})\Gamma(\frac{d-1}{2})} \Theta^{-\frac{1}{2}} (1-\Theta)^{\frac{d-3}{2}} & \text{if } 0 < \Theta < 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$\Theta \sim \text{Beta}\left(\frac{1}{2}, \frac{d-1}{2}\right).$$

$$(23)$$

This technique is used for the numerical tests in Section 5, and allows the efficient sampling of high dimensional spherically symmetric distributions where only a few parameters are of interest without generating all the remaining uninteresting parameters.

4 Estimating sampling errors in nested sampling parameter estimation

Following the discussion of sources of sampling errors in Section 3, we seek a method for correctly calculating parameter estimation sampling errors from a single nested sampling run. As no additional samples θ_i are available, a natural starting point is to utilise resampling techniques such as the jackknife (Tukey, 1958), bootstrap (Efron, 1979) and Bayesian bootstrap (Rubin, 1981), which estimate the uncertainty on inferences from a set of samples by calculating the variation when samples are re-weighted.

However, as described in Section 2.2, the uncertainty in nested sampling weights $w_i(\mathbf{t})$ produces additional sampling errors which are unique to the nested sampling process. These are not accounted for by naïvely applying jackknives and bootstraps to

posterior samples produced by nested sampling, and these approaches fail when tested numerically. We instead require a method for dividing runs in a manner that preserves the statistical properties of nested sampling. No such method exists in the literature, so we present one in the remainder of this section.

4.1 Dividing runs into threads

Skilling (2006) describes how several nested sampling runs r = 1, 2, ... with $n^{(r)}$ live points may be combined simply by merging the dead points and sorting by likelihood value. The combined sequence of dead points is equivalent to a single nested sampling run with $n = \sum_{r} n^{(r)}$ live points.

In fact, as we show now, the reverse procedure is also possible. A nested sampling run with n points can be unwoven into a set of n valid nested sampling runs, each with $n^{(r)} = 1$. We term these single live point runs *threads*. During nested sampling, each dead point i is replaced by a new point sampled uniformly within its iso-likelihood contour $\mathcal{L}(\theta) = \mathcal{L}_i$. Starting from each initial live point that is generated, one may follow this sequence of replacements down the set of dead points. This sub-sequence of dead points is in fact a nested sampling run with n = 1. More formally:

```
Result: n threads.
```

Data: Dead points and the iterations at which they were sampled for a nested sampling run with n live points.

Rank dead points by likelihood in ascending order; **while** $i \in n$ **do** make a new stack i; select one of the initial points sampled at the start of the run; move the point out to the stack i; **while** *iteration* < *final iteration* **do** | select point sampled at the iteration where previous point was replaced ("died"); | move the point to the stack i; **end end**

Algorithm 1: Splitting a nested sampling run into threads.

A few points are worthy of note:

- 1. splitting a run by randomly selecting some fraction of the dead points will not produce threads (i.e. single point nested sampling runs);
- 2. one may split a given nested sampling run into separate runs with $n^{(r)} \neq 1$ by first separating into threads, and then recombining threads as desired;
- 3. the algorithm can be easily adapted for varying numbers of live points by permitting it to select multiple points on contours where n increases. This can result

in constituent threads stopping or dividing into multiple threads part of the way through the run;

- 4. typically there is only one point which was sampled uniformly from the prior volume within each dead point *i*'s iso-likelihood contour $\mathcal{L}(\theta) = \mathcal{L}_i$ the point which replaced *i*. A sufficient condition for a nested sampling run to only have one unique division into threads is that $\mathcal{L}(X)$ is an injective function;
- 5. in order for the threads to be true nested sampling runs, care must be taken with the termination conditions conditions used. See Section D of the supplementary material for a full discussion.

Given that threads represent independent nested sampling runs, one may apply standard resampling techniques to the set of threads and approximate the entire sampling error distribution without making assumptions about its form. This works as the log X_i values of the dead points *i* from some run with *n* live points form a Poisson process with rate *n*, meaning the log X_j values of the dead points *j* of a single thread are a Poisson process of rate 1. For typical problems with computationally expensive likelihoods the computational cost of even a large number of resampling replications is negligible.

Having introduced a framework for applying resampling to nested sampling parameter estimation we now present an example method using bootstrap resampling.

4.2 Bootstrap estimate of sampling errors

Given *n* observations $\mathbf{x} = (x_1, \dots, x_n)$, the bootstrap (Efron, 1979) creates new data sets \mathbf{x}_b^* by drawing *n* samples from \mathbf{x} with replacement. This corresponds to approximating the probability distribution of a single data point *x* as

$$P(x) \approx \frac{1}{n} \sum_{i=1}^{n} \delta(x - x_i), \qquad (24)$$

where $\delta(x)$ is the Dirac delta function (Ivezić et al., 2014).

As the form of the distribution of sampling errors for a general nested sampling parameter estimation problem is not known, we use the non-parametric bootstrap. In this case the uncertainty on a quantity $T(\mathbf{x})$ calculated from the data can be estimated by calculating $T(\mathbf{x}_b^*)$ for a number of resampled data sets $b = 1, \ldots, B$. For example the bootstrap estimate of the standard error on $T(\mathbf{x})$ is

St.Dev.
$$[T(\mathbf{x})] = \sqrt{\frac{1}{B-1} \sum_{b=1}^{B} \left(T(\mathbf{x}_b^*) - \overline{T(\mathbf{x}_b^*)} \right)^2}, \text{ where } \overline{T(\mathbf{x}_b^*)} = \frac{1}{B} \sum_{b=1}^{B} T(\mathbf{x}_b^*).$$
(25)

There are many methods for calculating approximate credible intervals on $T(\mathbf{x})$ from bootstrap replications $\{T(\mathbf{x}_b^*)\}$ — see Efron and Tibshirani (1986) for a detailed

discussion. A simple approach from Johnson (2001) is to estimate the boundaries of the $100\alpha\%$ and $100(1-\alpha)\%$ credible regions⁵ as

C.I.<sub>100
$$\alpha\%$$</sub> $(T(\mathbf{x})) = 2T(\mathbf{x}) - G^{-1}(1-\alpha)$ (26)

C.I.<sub>100(1-
$$\alpha$$
)% ($T(\mathbf{x})$) = $2T(\mathbf{x}) - G^{-1}(\alpha)$, (27)</sub>

where $G^{-1}(x)$ is the inverse cumulative distribution of the bootstrap samples $\{T(\mathbf{x}_b^*)\}$. B = 50 is typically sufficient for an estimate of the standard deviation of a parameter estimate due to sampling errors, but depending on the method used credible intervals on parameter estimates may require 1,000 bootstrap replications or more (Efron and Tibshirani, 1986).

When the bootstrap is applied to nested sampling each observation x_i is a thread, and the number of observations is n. Calculating the quantity $T(\mathbf{x})$ involves first combining the set of threads \mathbf{x} into a single run using Skilling (2006)'s method (described in Section 4.1), then performing a standard nested sampling calculation including estimating the weight of each point $w_i(\mathbf{t})$ statistically. Including the same thread multiple times does not cause problems — repeated dead points $\theta_i = \theta_{i+1}$ are simply assigned the weights $w_i(\mathbf{t})$ and $w_{i+1}(\mathbf{t})$ respectively.

The following algorithm provides a set of bootstrap replications and an estimate of the standard deviation of sampling errors.

Result: Sampling errors and bootstrap replications for the nested sampling calculation T(dead points, weights).

Data: List of dead points and the steps they were sampled at. Divide dead points into a list of threads \mathbf{x} using Algorithm 1; while $b \in B$ do

create a list of *n* threads \mathbf{x}_b^* by sampling \mathbf{x} with replacement; calculate $T(\mathbf{x}_b^*) \equiv T(\text{dead points}_b^*, \text{weights}_b^*);$

 \mathbf{end}

calculate St.Dev.
$$[T(\mathbf{x})] = \sqrt{\frac{1}{B-1} \sum_{b=1}^{B} \left(T(\mathbf{x}_{b}^{*}) - \overline{T(\mathbf{x}_{b}^{*})} \right)^{2}}$$
.
Algorithm 2: Bootstrap sampling error calculation.

We find that bootstrap resampling gives better results than jackknife resampling, which fails to calculate sampling errors on credible intervals of posterior distributions of parameters such as $\text{C.I.}_{84\%}(\overline{\theta_1})$. The Bayesian bootstrap was not used as it gives each observation a non-integer weight, which requires modifying nested sampling's use of dead points to statistically estimate prior volume shrinkages.

Resampling techniques such as the bootstrap can generate many simulated runs with the same number of live points n as the original run. In comparison sampling error estimates from simply splitting a run into many smaller runs and assessing their variation perform poorly, as shown in Section C of the supplementary material.

⁵If the distribution of bootstrap replications $T(\mathbf{x}_{b}^{*})$ is skewed then the implied probability distribution of T is skewed in the opposite direction, as can be seen from (26) and (27). See Loredo (2012, Section 2) for a discussion.



Figure 4: Relative posterior mass as a function of $\log X \ (\propto \mathcal{L}(X)X)$ for Gaussian likelihoods (14) and Cauchy likelihoods (15) of different dimensions d with Gaussian priors (16). The lines are scaled so that the area under each of them is equal.

5 Numerical tests

Following Keeton (2011) we first test our new method using analytic cases where uncorrelated samples can be easily obtained from the prior within an iso-likelihood contour, allowing us to perform perfect nested sampling. This ensures our results are not affected by imperfect implementation of the nested sampling algorithm by a specific software.

As discussed in Section 3 perfect nested sampling parameter estimation problems depend on the likelihood $\mathcal{L}(\theta)$ and prior $\pi(\theta)$ only through the distribution of posterior mass $\mathcal{L}(X)$ and the distribution of parameters on iso-likelihood contours $P(f(\theta)|X)$, both of which are functions of both $\mathcal{L}(\theta)$ and $\pi(\theta)$. We therefore empirically test our method using a wide range of distributions of posterior mass, and examine several functions of parameters $f(\theta)$ in each case. We construct such tests using Gaussian likelihoods (14) and Cauchy likelihoods (15) of a variety of dimensions d, each with a Gaussian prior (16). The different distributions of posterior mass for different d are illustrated in Figure 4; the Cauchy distributions have extremely fat tails with significant sample weights throughout the range of log X values explored.

We use the termination conditions described by Handley et al. (2015b, Section 3.4), stopping when the estimated evidence contained in the live points is less than 10^{-4} times the evidence contained in dead points (see Section D of the supplementary material for a discussion of termination conditions for nested sampling parameter estimation). Numerical calculations for high-dimensional cases are performed in two dimensions using the technique described in Section 3.2.

As in Section 3 we denote the first component of the θ vector as $\theta_{\hat{1}}$, although by symmetry the results will be the same for any component. $\overline{\theta_{\hat{1}}}$ and $\overline{\theta_{\hat{1}}^2}$ are the first and second moments of the posterior distribution of $\theta_{\hat{1}}$, and the one-tailed Y% upper credible interval C.I._{Y%}($\overline{\theta_{\hat{1}}}$) is the value $\theta_{\hat{1}}^*$ for which $P(\theta_{\hat{1}} < \theta_{\hat{1}}^* | \mathcal{L}, \pi) = Y/100$. Figure 5: Sampling errors in a perfect nested sampling calculation for a 3dimensional Gaussian likelihood (14) and a uniform prior. The shading and black lines show the analytic posterior distribution and the 68% and 95% credible intervals. The red lines show the calculated posterior credible intervals for a nested sampling run with n = 100, and differ from the analytic answer due to sampling errors.



5.1 3-dimensional Gaussian example

We first test our bootstrap approach to estimating sampling errors on a 3-dimensional Gaussian likelihood (14) — Figure 5 illustrates sampling errors on posterior distributions of parameters θ in this case. Unlike the simulated weights method, the mean estimates of sampling errors from our method are very close to measurements of sampling errors from repeated calculations — this shown in the second row of Table 1. Furthermore the fractional variation of estimates from single runs around the mean estimate is similar to that from the simulated weights method, as shown in the fourth and fifth rows, indicating our method will give a reasonable estimate of sampling errors when only a single nested sampling run is available.

The final two rows of Table 1 show the empirical coverage rates for bootstrap credible intervals are very close to their nominal values. Figure 6 shows estimates of the full sampling error distribution from a single run nested sampling run using the bootstrap and simulated weights methods; the bootstrap results are much closer to the sampling errors observed in repeated calculations, and give accurate estimates of the 1σ and 2σ credible intervals.

Section \mathbf{E} of the supplementary material shows similar numerical tests for a 3dimensional Cauchy likelihood (15). Even for this challenging, fat-tailed distribution our method performs similarly to the Gaussian case, giving accurate mean error estimates and estimates of credible intervals with measured coverage similar to their nominal coverage.

	$\overline{ heta_{\hat{1}}}$	$\overline{ heta_{\hat{1}}^2}$	$\mathrm{C.I.}_{84\%}(\overline{\theta_{\hat{1}}})$
Repeated runs St.Dev.	0.032(0.2)	0.050(0.4)	0.055(0.4)
Bootstrap St.Dev. / Repeats St.Dev.	1.003(7)	0.998(7)	1.008(8)
Simulated w_i St.Dev. / Repeats St.Dev.	0.715(5)	0.882(6)	0.785(7)
Bootstrap St.Dev. estimate variation	7.5(1)%	8.6(1)%	17.7(3)%
Simulated w_i estimate variation	6.0(1)%	7.3(1)%	19.7(3)%
Bootstrap C.I. $_{95\%}$	0.053(3)	1.080(5)	1.077(7)
Bootstrap Mean ± 1 St.Dev. coverage	68.4%	68.2%	68.9%
Bootstrap C.I. $_{95\%}$ coverage	95.0%	93.4%	93.1%

Table 1: Sampling errors for a 3-dimensional Gaussian likelihood (14), a Gaussian prior (16) and n = 200. The first row shows the standard deviation of 10,000 nested sampling calculations. The second and third rows show the mean of 2,000 error estimates from the bootstrap and simulated weights methods respectively as a ratio to the error observed from repeated calculations; 200 weight simulations and 200 bootstrap replications were used for each run. The fourth and fifth rows show the standard deviations of sampling error estimates for both methods as a percentage of the mean estimate. The sixth row shows the mean of 100 bootstrap estimates of the one-tailed 95% credible interval on the calculation result given the sampling error, each using 1,000 bootstrap replications. The final two rows show the empirical coverage of the bootstrap standard error and 95% credible interval from the 10,000 repeated calculations. Numbers in brackets show the error on the final digit.

5.2 Sampling errors in different dimensions

We now verify the bootstrap method's accuracy for Gaussian (14) and Cauchy (15) likelihoods of between 2 and 50 dimensions. Figure 7 shows bootstrap sampling error estimates accurately match the errors measured from repeated calculations, even for the challenging fat-tailed Cauchy distribution. In contrast the simulated weights method consistently underestimates the sampling errors in parameter estimation, although as expected it is accurate for errors on the evidence log \mathcal{Z} . See Section B of the supplementary material for a detailed discussion of the simulated weights method.

As the dimension d increases, Figure 7 shows parameter estimation errors decreasing and the evidence errors increasing (with a constant number of live points n). This effect is due to the posterior being contained in a smaller fraction of the prior volume in higher dimensions. In the spherically symmetric cases considered, the range of log Xto be explored increases approximately linearly with the dimension d, as can be seen in Figure 4. With a constant number of live points, the number of samples is therefore also approximately proportional to d.

In parameter estimation from posterior samples only points' relative weights matter, so the increased number of samples in higher dimension problems typically increases accuracy as can be seen in Figures 7(a) and 7(b). However for high dimensional Cauchy



Figure 6: Estimated distributions of sampling errors for parameter estimation with a 3-dimensional Gaussian likelihood (14) for perfect nested sampling with n = 200. For each estimator the first plot uses values from 5,000 nested sampling runs; the second and third plot are calculated from a single nested sampling run and use 5,000 simulated weights and bootstrap replications. The bootstrap distributions are calculated using (26) and (27). The simulated weights and bootstrap values were adjusted by subtracting the difference between their run's expected value for each estimator and its analytical value to line up the distributions. The colour scale shows the fraction of the cumulative probability distribution lying between some region and the median.

likelihoods (15) the posterior mass is spread over a wide range of $\log X$ values, so errors in the relative weights of points become large in high dimensions⁶.

For $\log \mathcal{Z}$ the dominant error is in the absolute value of point weights, which is approximately proportional to the square root of the number of steps required to reach the posterior (Skilling, 2006). $\log \mathcal{Z}$ errors are therefore approximately proportional to \sqrt{d} when n is constant, as can be seen in Figure 7(c).

6 Application to existing nested sampling software

Nested sampling software such as MULTINEST and POLYCHORD can be easily modified to output information about the step at which dead points were sampled and give sampling error estimates using bootstrap resampling of threads. We are currently working on incorporating this into future releases of POLYCHORD.

⁶When the errors in points' relative weights become dominant the simulated weights method captures the majority of the sampling error, as can be seen for high dimensional Cauchy distributions in Figure 7(b)





Figure 7: Sampling errors for likelihoods of different dimensions d; all use use ddimensional Gaussian priors (16) and n = 100. Solid lines show the standard deviation of the results of 2,000 calculations. Dashed and dotted lines show the mean of 500 standard error estimates using the bootstrap and simulated weights methods respectively.

Sampling error estimates from our approach will be accurate provided the software is performing nested sampling approximately correctly. However such software can only approximately sample randomly from the prior within iso-likelihood contours — this may result in additional errors which are specific to a given implementation and which may not be captured by general methods such as resampling threads. Tuning parameters such as "num_repeats" (the number of slice samples taken between dead points) in POLYCHORD allow reduced correlation between samples at a higher computational cost. Testing sampling error estimates from our method against those from repeated calculations can be used to detect implementation-specific errors and to select appropriate values for tuning parameters.

In addition our algorithm for dividing nested sampling runs could be used to detect implementation-specific errors by testing the difference in correlation between threads from the same run and from different runs. In principle this could allow estimates of sampling errors to be corrected for the effects of correlations between threads.

We now demonstrate our method's application to nested sampling results produced with POLYCHORD.

6.1 Sampling errors on data fitting with POLYCHORD

We fit a set of points $\mathcal{D} = \{x_i, y_i\}$ with normally distributed errors σ_y on the y values using a sinusoid

$$y(x) = A\sin(\omega x + \phi). \tag{28}$$

The likelihood is then

$$\mathcal{L}(\theta) = \prod_{i} \frac{1}{\sqrt{2\pi\sigma_y^2}} e^{-(y_i - y(x))^2/2\sigma_y^2},$$
(29)

where $\theta = (A, \omega, \phi)$ and we use a uniform prior for $A \in (0, 1)$, $\omega \in (0, 10)$ and $\phi \in (-\pi/2, \pi/2)$. Numerical tests use 40 data points sampled from $y(x) = \frac{1}{2}\sin(2\pi x)$ with Gaussian noise of size $\sigma_y = 0.2$ added to the y values; y(x) and the data points are shown in Figure 8 and the posterior distribution of y(x) given the data is shown in Figure 9. Posterior distributions on A, ω and ϕ can be calculated with nested sampling — these are illustrated in Figure 10 along with example sampling errors.

Table 2 shows sampling errors from POLYCHORD with num_repeats = 15 — the default value for a 3-dimensional problem. As in perfect nested sampling, our bootstrap estimates of the standard error agree with the variation in results observed, and the observed coverage of credible intervals is close to their nominal coverage. This implies that num_repeats = 15 is sufficient for POLYCHORD to perform parameter estimation accurately in this case.

7 Conclusion

Sampling errors in nested sampling parameter estimation arise principally from two sources: uncertain sample weights $w_i(\mathbf{t})$, and approximating the average of a function



Figure 8: The true y(x) function and the 40 data points used in numerical tests. Data has Gaussian noise of size $\sigma_y = 0.2$.



Figure 9: Posterior distribution on y(x) given the data \mathcal{D} . The color scale indicates credible intervals.



Figure 10: Posterior distributions and sampling errors from fitting a sinusoid to data (29) using POLY-CHORD. The shading and black lines show an accurate calculation of posterior distribution and the 68% and 95% credible intervals from combining 1,000 nested sampling runs with n= 100. The red lines show the calculated posterior credible intervals for a single nested sampling run with n = 100, and differ from the true posterior due to sampling errors.

of parameters on each iso-likelihood contour $\tilde{f}(X_i)$ with a single sample $f(\theta_i)$. The latter error is not present in evidence calculation and has been previously ignored. The added stochasticity from sampling each iso-likelihood contour makes nested sampling parameter estimation a 2-dimensional problem, with a dependence on both the distribution of posterior mass $\mathcal{L}(X)$ and the distribution of parameter values $P(f(\theta)|X)$ on each iso-likelihood contour. We proposed a new diagram for representing both aspects of the calculation, and presented it in Figures 3(a) to 3(e).

	Α	ω	$\mathrm{C.I.}_{84\%}(\omega)$
Repeated runs St.Dev.	$0.247(6) \cdot 10^{-2}$	0.012(0.3)	0.019(0.4)
Bootstrap St.Dev. / Repeats St.Dev.	0.98(2)	0.98(2)	0.97(2)
Simulated w_i St.Dev. / Repeats St.Dev.	0.71(2)	0.71(2)	0.73(2)
Bootstrap St.Dev. estimate variation	8.8(2)%	9.3(2)%	19.1(4)%
Simulated w_i estimate variation	6.4(1)%	6.6(1)%	21.5(5)%
Bootstrap C.I. $_{95\%}$	0.577(0.2)	6.370(0.8)	6.632(1)
Bootstrap Mean ± 1 St.Dev. coverage	69.4%	69.3%	66.9%
Bootstrap C.I. $_{95\%}$ coverage	93.4%	95.0%	95.1%

Table 2: Sampling errors for the sinusoid fitting likelihood (29) using POLYCHORD with n = 100. The first row shows the standard deviation of 1,000 nested sampling calculations. The second and third rows show the mean of 1,000 error estimates from the bootstrap and simulated weights methods respectively as a ratio to the error observed from repeated calculations; 200 weight simulations and 200 bootstrap replications were used for each run. The fourth and fifth rows show the standard deviations of sampling error estimates for both methods as a percentage of the mean estimate. The sixth row shows the mean of 100 bootstrap estimates of the one-tailed 95% credible interval on the calculation result, each using 1,000 bootstrap replications. The final two rows show the empirical coverage of the bootstrap standard error and 95% credible interval from the 1,000 repeated calculations. Numbers in brackets show the error on the final digit.

Estimating sampling errors is vital for interpreting the results of a nested sampling calculation, as well as for allocating computational resources — for example by choosing an appropriate number of live points. However the current approach (the simulated weights method) underestimates sampling errors as it does not account for approximating $\tilde{f}(X_i)$ with a single sample $f(\theta_i)$. We proposed a new method for estimating sampling errors using our new algorithm (Algorithm 1) for dividing a nested sampling run into single live point runs ("threads"), which can then be resampled with techniques such as the bootstrap. This works as the log X_i values of the dead points *i* from some nested sampling run with *n* live points form a Poisson process with rate *n*, meaning the log X_j values of the dead points *j* of a single thread are a Poisson process of rate 1.

Our method shows accurate and robust estimation of sampling errors in parameter estimation in empirical tests, and compares favourably to the other methods discussed. The new method can be easily incorporated into existing nested sampling software, and will be reliable provided the implementation is performing the nested sampling algorithm accurately. We are currently working on including nested sampling run division and sampling error estimates from our method in future versions of POLYCHORD.

References

- Bezáková, I., Stefankovic, D., Vazirani, V. V., and Vigoda, E. (2008). "Accelerating simulated annealing for the permanent and combinatorial counting problems." SIAM, 37(5): 1429–1454. 3
- Chopin, N. and Robert, C. P. (2010). "Properties of nested sampling." *Biometrika*, 97(3): 741–755. 3, 5, 6
- Del Moral, P., Doucet, A., and Jasra, A. (2006). "Sequential Monte Carlo samplers." Journal of the Royal Statistical Society. Series B: Statistical Methodology, 68(3): 411– 436. 3
- Doss, C., Flegal, J., Jones, G., and Neath, R. (2015). "Markov chain monte carlo estimation of quantiles." *Electronic Journal of Statistics*, 8: 2448–2478. 3
- Efron, B. (1979). "Bootstrap Methods: Another Look at the Jackknife." *The Annals of Statistics*, 7(1): 1–26. 11, 13
- Efron, B. and Tibshirani, R. (1986). "Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy." *Statistical Science*, 1(1): 54–77. 13, 14
- Feroz, F. and Hobson, M. P. (2008). "Multimodal nested sampling: An efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses." Monthly Notices of the Royal Astronomical Society, 384(2): 449–463. 1
- Feroz, F., Hobson, M. P., and Bridges, M. (2009). "MultiNest: An efficient and robust Bayesian inference tool for cosmology and particle physics." *Monthly Notices of the Royal Astronomical Society*, 398(4): 1601–1614. 1
- Feroz, F., Hobson, M. P., Cameron, E., and Pettitt, A. N. (2013). "Importance Nested Sampling and the MultiNest Algorithm." arXiv preprint arXiv:1306.2144, 28. 1, 3
- Flegal, J. M., Haran, M., and Jones, G. L. (2008). "Markov Chain Monte Carlo: Can We Trust the Third Significant Figure?" *Statistical Science*, 23(2): 250–260. 3
- Handley, W. J., Hobson, M. P., and Lasenby, A. N. (2015a). "PolyChord: Nested sampling for cosmology." Monthly Notices of the Royal Astronomical Society: Letters, 450(1): L61–L65. 1
- (2015b). "PolyChord: next-generation nested sampling." Monthly Notices of the Royal Astronomical Society, 15: 1–15. 1, 5, 15, 31
- Huber, M. and Schott, S. (2014). "Random construction of interpolating sets for highdimensional integration." Journal of Applied Probability, 51(1): 92–105. 3
- Ivezić, Ż., Connolly, A., VanderPlas, J., and Gray, A. (2014). Statistics, Data Mining, and Machine Learning in Astronomy. Princeton University Press. 13
- Johnson, R. W. (2001). "An Introduction to the Bootstrap." *Teaching Statistics*, 23(2): 49–54. 14

- Keeton, C. R. (2011). "On statistical uncertainty in nested sampling." Monthly Notices of the Royal Astronomical Society, 414(2): 1418–1426. 2, 5, 15, 29
- Liu, J., Nordman, D. J., and Meeker, W. Q. (2016). "The Number of MCMC Draws Needed to Compute Bayesian Credible Bounds." *The American Statistician*, 06340: 1–27. 3
- Loredo, T. J. (2012). "Bayesian astrostatistics: a backward look to the future." Astrostatistical Challenges for the New Astronomy, 15–40. 14
- MacKay, D. J. C. (2003). Information Theory, Inference, and Learning Algorithms. Cambridge University Press. 1
- Planck Collaboration (2016). "Planck 2015. XX. Constraints on inflation." Astronomy & Astrophysics, 594: A20. 1
- Rubin, D. B. (1981). "The Bayesian Bootstrap." Annals of Statistics, 9(1): 130–134. 11, 25
- Skilling, J. (2006). "Nested sampling for general Bayesian computation." Bayesian Analysis, 1(4): 833–860. 1, 3, 4, 5, 12, 14, 18, 30
- Tukey, J. W. (1958). "Bias and Confidence in Not-Quite Large Samples." The Annals of Mathematical Statistics, 29: 614. 11

Acknowledgments

The authors are grateful to the referees and editors for their helpful comments and suggestions, which have greatly improved the paper. In particular we thank one of the reviewers for providing the justification of resampling nested sampling threads in terms of Poisson processes.