

---

# Astronomical Data Analysis

Alan Heavens

Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, afh@roe.ac.uk  
Expanded version of lecture at STFC Introductory School, University of Glasgow, August 2011

## 1 Introduction

In the lectures I briefly covered a number of topics in astronomical data analysis. These notes expand on the lecture, concentrating on general techniques which are common in astronomy. In fact they have very general applicability, for problems in which the data are interpreted in the context of a theoretical model, and thus lend themselves to a Bayesian treatment.

We consider the general problem of estimating parameters from data, and consider how one can use Fisher matrices to analyse survey designs before any data are taken, to see whether the survey will actually do what is required. We outline numerical methods for estimating parameters from data, including Monte Carlo Markov Chains and the Hamiltonian Monte Carlo method. We also look at Model Selection, which covers various scenarios such as whether an extra parameter is preferred by the data, or answering wider questions such as which theoretical framework is favoured, using General Relativity and braneworld gravity as an example. These notes are not a literature review, so there are relatively few references. A slightly expanded version of these notes is available [10]. For further reading, [22] is a textbook introduction to Bayesian methods, and [6] is a good resource for MCMC.

After this introduction, the sections are:

- Parameter Estimation
- Fisher Matrix analysis
- Numerical methods for Parameter Estimation
- Model Selection

### 1.1 Notations:

- *Data* will be called  $\mathbf{x}$ , or  $x$ , or  $x_i$ , and is written as a vector, even if it is a 2D image.
- *Model parameters* will be called  $\boldsymbol{\theta}$ , or  $\theta$  or  $\theta_\alpha$

### 1.2 Inverse problems

Most data analysis problems are *inverse problems*. You have a set of data  $\mathbf{x}$ , and you wish to interpret the data in some way. Typical classifications are:

- Hypothesis testing
- Parameter estimation
- Model selection

Cosmological examples of the first type include

- Are CMB data consistent with the hypothesis that the initial fluctuations were gaussian, as predicted (more-or-less) by the simplest inflation theories?

- Are large-scale structure observations consistent with the hypothesis that the Universe is spatially flat?

Cosmological examples of the second type include

- In the Big Bang model, what is the value of the matter density parameter?
- What is the value of the Hubble constant?

Model selection can include slightly different types (but are mostly concerned with larger, often more qualitative questions):

- Do cosmological data favour the Big Bang theory or the Steady State theory?
- Is the gravity law General Relativity or higher-dimensional?
- Is there evidence for a non-flat Universe?

Note that the notion of a *model* can be a completely different paradigm (the first two examples), or basically the same model, but with a different parameter set. In the third example, we are comparing a flat Big Bang model with a non-flat one. The latter has an additional parameter, and is considered to be a different model. Similar problems exist elsewhere in astrophysics, such as how many absorption-line systems are required to account for an observed feature?

These lectures will principally be concerned with questions of the last two types, *parameter estimation* and *model selection*, but will also touch on experimental design and error forecasting. Hypothesis testing can be treated in a similar manner.

## 2 Parameter estimation

We collect some data, and wish to interpret them in terms of a *model*. A model is a theoretical framework which we assume is true. It will typically have some parameters  $\theta$  in it, which you want to determine. The goal of parameter estimation is to provide estimates of the parameters, and their errors, or ideally the whole probability distribution of  $\theta$ , given the data  $\mathbf{x}$ . This is called the *posterior probability distribution* i.e. it is the probability that the parameter takes certain values, *after* doing the experiment (as well as assuming a lot of other things):

$$p(\theta|\mathbf{x}). \quad (1)$$

This is an example of RULE 1<sup>1</sup>: Start by thinking about what it is you want to know, and write it down mathematically.

From  $p(\theta|\mathbf{x})$  one can calculate the expectation values of the parameters, and their errors. Note that we are immediately taking a Bayesian view of probability, as a *degree of belief*, rather than a frequency of occurrence in a set of trials.

### 2.1 Forward modelling

Often, what may be easily calculable is not this, rather the opposite,  $p(\mathbf{x}|\theta)^2$ . The opposite is sometimes referred to as *forward modelling* - i.e. if we know what the parameters are, we can compute the expected distribution of the data. Examples of forward modelling distributions include the common ones - binomial, Poisson, gaussian etc. or may be more complex, such as the predictions for the CMB power spectrum as a function of cosmological parameters. As a concrete example, consider a model which is a gaussian with mean  $\mu$  and variance  $\sigma^2$ . The model has two parameters,  $\theta = (\mu, \sigma)$ , and the probability of a single variable  $x$  given the parameters is

$$p(x|\theta) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right], \quad (2)$$

but this is not what we actually want. However, we can relate this to  $p(\theta|x)$  using Bayes' Theorem, here written for a more general data vector  $\mathbf{x}$ :

$$p(\theta|\mathbf{x}) = \frac{p(\theta, \mathbf{x})}{p(\mathbf{x})} = \frac{p(\mathbf{x}|\theta)p(\theta)}{p(\mathbf{x})}. \quad (3)$$

<sup>1</sup> There is no rule  $n: n > 1$ .

<sup>2</sup> If you are confused about  $p(A|B)$  and  $p(B|A)$  consider if A=pregnant and B=female.  $p(A|B)$  is a few percent,  $p(B|A)$  is unity

- $p(\boldsymbol{\theta}|\mathbf{x})$  is the *posterior* probability for the parameters.
- $p(\mathbf{x}|\boldsymbol{\theta})$  is called the *Likelihood* and given its own symbol  $L(\mathbf{x}; \boldsymbol{\theta})$ .
- $p(\boldsymbol{\theta})$  is called the *prior*, and expresses what we know about the parameters prior to the experiment being done. This may be the result of previous experiments, or theory (e.g. some parameters, such as the age of the Universe, may have to be positive). In the absence of any previous information, the prior is often assumed to be a constant (a ‘flat prior’).
- $p(\mathbf{x})$  is the *evidence*.

For parameter estimation, the evidence simply acts to normalise the probabilities,

$$p(\mathbf{x}) = \int d\boldsymbol{\theta} p(\mathbf{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta}) \quad (4)$$

and the *relative* probabilities of the parameters do not depend on it, so it is often ignored and not even calculated.

However, the evidence does play an important role in *model selection*, when more than one theoretical model is being considered, and one wants to choose which model is most likely, whatever the parameters are. We turn to this later.

Actually all the probabilities above should be conditional probabilities, given any prior information  $I$  which we may have. For clarity, I have omitted these for now.  $I$  may be the result of previous experiments, or may be a theoretical prior, in the absence of any data. In such cases, it is common to adopt the *principle of indifference* and assume that all values of the parameter(s) is (are) equally likely, and take  $p(\boldsymbol{\theta}) = \text{constant}$  (perhaps within some finite bounds, or if infinite bounds, set it to some arbitrary constant and work with an unnormalised prior). This is referred to as a *flat prior*. Other choices can sometimes be justified, e.g. a uniform distribution of  $\ln \theta$ , if  $\theta$  is a parameter which scales the expectation value of the data, and one has no prior idea of which decade it lies in.

Thus for flat priors, we have simply

$$p(\boldsymbol{\theta}|\mathbf{x}) \propto L(\mathbf{x}; \boldsymbol{\theta}). \quad (5)$$

Although we may have the full probability distribution for the parameters, often one simply uses the peak of the distribution as the estimate of the parameters. This is then a *Maximum Likelihood* estimate. Note that if the priors are not flat, the peak in the posterior  $p(\boldsymbol{\theta}|\mathbf{x})$  is not necessarily the maximum likelihood estimate.

A ‘rule of thumb’ is that if the priors are assigned theoretically, and they influence the result significantly, the data are usually not good enough. (If the priors come from previous experiment, the situation is different - we can be more certain that we really have some prior knowledge in this case).

Finally, note that this method does not generally give a goodness-of-fit, only relative probabilities. It is still common to compute  $\chi^2$  at this point to check the fit is sensible.

## 2.2 Errors

Let us assume we have a posterior probability distribution, which is single-peaked. Two common estimators (indicated by a hat:  $\hat{\boldsymbol{\theta}}$ ) of the parameters are the peak (most probable) values, or the mean,

$$\hat{\boldsymbol{\theta}} = \int d\boldsymbol{\theta} \boldsymbol{\theta} p(\boldsymbol{\theta}|\mathbf{x}). \quad (6)$$

An estimator is *unbiased* if its expectation value is the true value  $\boldsymbol{\theta}_0$ :

$$\langle \hat{\boldsymbol{\theta}} \rangle = \boldsymbol{\theta}_0. \quad (7)$$

Let us assume for now that the prior is flat, so the posterior is proportional to the likelihood. This can be relaxed. Close to the peak, a Taylor expansion of the log likelihood implies that locally it is a multivariate gaussian *in parameter space*:

$$\ln L(\mathbf{x}; \boldsymbol{\theta}) = \ln L(\mathbf{x}; \boldsymbol{\theta}_0) + \frac{1}{2}(\boldsymbol{\theta}_\alpha - \boldsymbol{\theta}_{0\alpha}) \frac{\partial^2 \ln L}{\partial \boldsymbol{\theta}_\alpha \partial \boldsymbol{\theta}_\beta} (\boldsymbol{\theta}_\beta - \boldsymbol{\theta}_{0\beta}) + \dots \quad (8)$$

or

$$L(\mathbf{x}; \boldsymbol{\theta}) = L(\mathbf{x}; \boldsymbol{\theta}_0) \exp \left[ -\frac{1}{2}(\boldsymbol{\theta}_\alpha - \boldsymbol{\theta}_{0\alpha}) H_{\alpha\beta} (\boldsymbol{\theta}_\beta - \boldsymbol{\theta}_{0\beta}) \right]. \quad (9)$$

The Hessian matrix  $H_{\alpha\beta} \equiv -\frac{\partial^2 \ln L}{\partial \boldsymbol{\theta}_\alpha \partial \boldsymbol{\theta}_\beta}$  controls whether the estimates of  $\theta_\alpha$  and  $\theta_\beta$  are correlated or not. If it is diagonal, the estimates are uncorrelated. Note that this is a statement about *estimates* of the quantities, not the

quantities themselves, which may be entirely independent, but if they have a similar effect on the data, their estimates may be correlated. Note that in cases of practical interest, the likelihood may not be well described by a multivariate gaussian at levels which set the interesting credibility<sup>3</sup> levels (e.g. 68%). We turn later to how to proceed in such cases.

### 2.3 Conditional and marginal errors

If we fix all the parameters except one, then the error is given by the curvature along a line through the likelihood (posterior, if prior is not flat):

$$\sigma_{\text{conditional},\alpha} = \frac{1}{\sqrt{H_{\alpha\alpha}}}. \quad (10)$$

This is called the *conditional error*, and is the minimum error bar attainable on  $\theta_\alpha$  if all the other parameters are known. *It is rarely relevant and should almost never be quoted.*

### 2.4 Marginalising over a gaussian likelihood

The marginal distribution of  $\theta_1$  is obtained by integrating over the other parameters:

$$p(\theta_1) = \int d\theta_2 \dots d\theta_N p(\theta) \quad (11)$$

a process which is called *marginalisation*. Often one sees marginal distributions of all parameters in pairs, as a way to present some complex results. In this case two variables are left out of the integration.

If you plot such error ellipses, you *must* say what contours you plot. If you say you plot  $1\sigma$  and  $2\sigma$  contours, I don't know whether this is for the joint distribution (i.e. 68% of the probability lies within the inner contour), or whether 68% of the probability of a single parameter lies within the bounds projected onto a parameter axis. The latter is a  $1\sigma$ , single-parameter error contour (and corresponds to  $\Delta\chi^2 = 1$ ), whereas the former is a  $1\sigma$  contour for the joint distribution, and corresponds to  $\Delta\chi^2 = 2.3$ .

Note that  $\Delta\chi^2 = \chi^2 - \chi^2(\text{minimum})$ , where

$$\chi^2 = \sum_i \frac{(x_i - \mu_i)^2}{\sigma_i^2} \quad (12)$$

for data  $x_i$  with  $\mu_i = \langle x_i \rangle$  and variance  $\sigma_i^2$ . If the data are correlated, this generalises to

$$\chi^2 = \sum_{ij} (x_i - \mu_i) C_{ij}^{-1} (x_j - \mu_j) \quad (13)$$

where  $C_{ij} = \langle (x_i - \mu_i)(x_j - \mu_j) \rangle$ .

For other dimensions, see Table 1, or read...

### The Numerical Recipes bible, chapter 15.6 [20]

Read it. Then, when you need to plot some error contours, read it again.

**Table 1.**  $\Delta\chi^2$  for joint parameter estimation for 1, 2 and 3 parameters.

| $\sigma$  | p      | M=1  | M=2  | M=3  |
|-----------|--------|------|------|------|
| $1\sigma$ | 68.3%  | 1.00 | 2.30 | 3.53 |
| $2\sigma$ | 95.4%  | 4.00 | 6.17 | 8.02 |
| $3\sigma$ | 99.73% | 9.00 | 11.8 | 14.2 |

Note that some of the results I give assume the likelihood (or posterior) is well-approximated by a multivariate gaussian. This may not be so. If your posterior is a single peak, but is not well-approximated by a multivariate

<sup>3</sup> Often this is called the confidence level, but the credibility level is the appropriate terminology for Bayesian statistics. The 95% confidence interval is a frequentist construct: it is the interval which, for a given model and parameter, then, allowing for all data which MIGHT have been obtained, the procedure for constructing the confidence interval would include the true parameter 95% of the time.

gaussian, label your contours with the enclosed probability. If the likelihood is complicated (e.g. multimodal), then you may have to plot it and leave it at that - reducing it to a maximum likelihood point and error matrix is not very helpful. Not that in this case, the mean of the posterior may be unhelpful - it may lie in a region of parameter space with a very small posterior.

A multivariate gaussian likelihood is a common assumption, so it is useful to compute marginal errors for this rather general situation. The simple result is that the marginal error on parameter  $\theta_\alpha$  is

$$\sigma_\alpha = \sqrt{(\mathbf{H}^{-1})_{\alpha\alpha}}. \quad (14)$$

Note that we invert the Hessian matrix, and then take the square root of the diagonal components. Let us prove this important result. In practice it is often used to estimate errors for a future experiment, where we deal with the expectation value of the Hessian, called the *Fisher Matrix*:

$$\mathbf{F}_{\alpha\beta} \equiv \langle \mathbf{H}_{\alpha\beta} \rangle = \left\langle -\frac{\partial^2 \ln L}{\partial \theta_\alpha \partial \theta_\beta} \right\rangle. \quad (15)$$

We will have much more to say about Fisher matrices later. The expected error on  $\theta_\alpha$  is thus

$$\sigma_\alpha = \sqrt{(\mathbf{F}^{-1})_{\alpha\alpha}}. \quad (16)$$

It is always at least as large as the expected conditional error. Note: this result applies for gaussian-shaped likelihoods, and is useful for experimental design. For real data, you would do the marginalisation a different way - see later.

To prove this, we can use characteristic functions. See expanded version of these notes[10]

### 3 Fisher Matrix Analysis

This has been adapted from [24] (hereafter TTH).

How accurately can we estimate model parameters from a given data set? This question was basically answered 60 years ago [5], and we will now summarize the results, which are both simple and useful.

Suppose for definiteness that our data set consists of  $N$  real numbers  $x_1, x_2, \dots, x_N$ , which we arrange in an  $N$ -dimensional vector  $\mathbf{x}$ . These numbers could for instance denote the measured temperatures in the  $N$  pixels of a CMB sky map, the counts-in-cells of a galaxy redshift survey,  $N$  coefficients of a Fourier expansion of an observed galaxy density field, or the number of gamma-ray bursts observed in  $N$  different flux bins. Before collecting the data, we think of  $\mathbf{x}$  as a random variable with some probability distribution  $L(\mathbf{x}; \boldsymbol{\theta})$ , which depends in some known way on a vector of  $M$  model parameters  $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_M)$ .

Such model parameters might for instance be the spectral index of density fluctuations, the Hubble constant  $h$ , the cosmic density parameter  $\Omega$  or the mean redshift of gamma-ray bursts. We will let  $\boldsymbol{\theta}_0$  denote the true parameter values and let  $\boldsymbol{\theta}$  refer to our estimate of  $\boldsymbol{\theta}$ . Since  $\boldsymbol{\theta}$  is some function of the data vector  $\mathbf{x}$ , it too is a random variable. For it to be a good estimate, we would of course like it to be unbiased, *i.e.*,

$$\langle \boldsymbol{\theta} \rangle = \boldsymbol{\theta}_0, \quad (17)$$

and give as small error bars as possible, *i.e.*, minimize the standard deviations

$$\Delta\theta_\alpha \equiv (\langle \theta_\alpha^2 \rangle - \langle \theta_\alpha \rangle^2)^{1/2}. \quad (18)$$

In statistics jargon, we want the BUE  $\theta_\alpha$ , which stands for the ‘‘Best Unbiased Estimator’’.

A key quantity in this context is the so-called *Fisher information matrix*, defined as

$$\mathbf{F}_{\alpha\beta} \equiv \left\langle \frac{\partial^2 \mathcal{L}}{\partial \theta_\alpha \partial \theta_\beta} \right\rangle, \quad (19)$$

where

$$\mathcal{L} \equiv -\ln L. \quad (20)$$

Another key quantity is the *maximum likelihood estimator*, or *ML-estimator* for brevity, defined as the parameter vector  $\boldsymbol{\theta}_{\text{ML}}$  that maximizes the likelihood function  $L(\mathbf{x}; \boldsymbol{\theta})$ .

Using this notation, a number of powerful theorems have been proven (see *e.g.* [14],[13]):

1. For any unbiased estimator,  $\Delta\theta_\alpha \geq 1/\sqrt{F_{\alpha\alpha}}$  (the *Cramér-Rao* inequality).
2. If an unbiased estimator attaining (“saturating”) the Cramér-Rao bound exists, it is the ML estimator (or a function thereof).
3. The ML-estimator is asymptotically BUE.

The first of these theorems thus places a firm lower limit on the error bars that one can attain, regardless of which method one is using to estimate the parameters from the data. You won’t do better, but you might do worse.

The normal case is that the other parameters are estimated from the data as well, in which case, as we have seen, the minimum standard deviation rises to

$$\Delta\theta_\alpha \geq (F^{-1})_{\alpha\alpha}^{1/2}. \quad (21)$$


This is called the *marginal error*, and I reemphasise that this is normally the relevant error to quote.

The second theorem shows that maximum-likelihood (ML) estimates have quite a special status: if there is a best method, then the ML-method is the one. Finally, the third result basically tells us that in the limit of a very large data set, the ML-estimate for all practical purposes is the best estimate, the one that for which the Cramér-Rao inequality becomes an equality<sup>4</sup>. It is these nice properties that have made ML-estimators so popular.

Note that conditional and marginal errors coincide if  $F$  is diagonal. If it is not, then the *estimates* of the parameters are correlated (even if the parameters themselves are uncorrelated). e.g. in the example shown in Fig. 1, estimates of the baryon density parameter  $\Omega_b$  and the dark energy equation of state  $w \equiv p/\rho c^2$  are strongly correlated with WMAP data alone, so we will tend to overestimate both  $\Omega_b$  and  $w$ , or underestimate both. However, the value of  $\Omega_b$  in the Universe has nothing obvious to do with  $w$  - they are independent.

Note that it is quite common to see plots like this, where all parameters except for two have been marginalised over. To produce such plots, a  $2 \times 2$  matrix ( $F^{(2)}$ ) of relevant rows and columns (corresponding to the two parameters whose marginal probability distribution is to be plotted) is extracted from  $F^{-1}$ . What is plotted is then

$$\ln L = \ln L_0 - \frac{1}{2}(\theta - \theta_0)_\alpha F_{\alpha\beta}^{(2)}(\theta - \theta_0)_\beta \quad \alpha, \beta = 1, 2.$$



**Fig. 1.** The expected error ellipses for cosmological parameters ( $\sigma_8$ , baryon density parameter  $\Omega_b$ , and dark energy equation of state  $w \equiv p/\rho c^2$ ) from a 3D weak lensing survey of 1000 square degrees, with a median redshift of 1 and a photometric redshift error of 0.15. Probabilities are marginalised over all other parameters, except that  $n = 1$  and a flat Universe are assumed. Dark ellipses represent a prior from WMAP, pale represents the 3D lensing survey alone, and the central ellipses show the combination (from Kitching T., priv. comm.).

<sup>4</sup> This is sometimes called ‘saturating the Cramér-Rao bound’

## Combining experiments

If the experiments are independent, you can simply add the Fisher matrices (why?). Note that the marginal error ellipses (marginalising over all but two variables) in the combined dataset can be much smaller than you might expect, given the marginal error ellipses for the individual experiments, because the operations of adding the experimental data and marginalising do not commute.

### 3.1 The Gaussian Case

Let us now explicitly compute the Fisher information matrix for the case when the probability distribution is Gaussian, *i.e.*, where (dropping an irrelevant additive constant  $N \ln[2\pi]$ )

$$2\mathcal{L} = \ln \det C + (\mathbf{x} - \boldsymbol{\mu})C^{-1}(\mathbf{x} - \boldsymbol{\mu})^T, \quad (22)$$

where in general both the mean vector  $\boldsymbol{\mu} \equiv \langle \mathbf{x} \rangle$  and the covariance matrix

$$C = \langle (\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \rangle \quad (23)$$

depend on the model parameters  $\boldsymbol{\theta}$ . Although vastly simpler than the most general situation, the Gaussian case is nonetheless general enough to be applicable to a wide variety of problems in cosmology. Some interesting mathematics allows us to compute the Fisher matrix, which is simply the expectation value of  $-\ln L$  (see [24, 10] for more details), giving the important result that

$$F_{\alpha\beta} = \langle \mathcal{L}_{,\alpha\beta} \rangle = \frac{1}{2} \text{Tr}[C^{-1}C_{,\alpha}C^{-1}C_{,\beta} + C^{-1}M_{\alpha\beta}], \quad (24)$$

where we have defined the matrix  $M_{\alpha\beta} \equiv \langle D_{,\alpha\beta} \rangle = \boldsymbol{\mu}_{,\alpha} \boldsymbol{\mu}_{,\beta}^T + \boldsymbol{\mu}_{,\beta} \boldsymbol{\mu}_{,\alpha}^T$ , and  $C_{,\alpha} \equiv \partial C / \partial \theta_\alpha$ .

Note that *the Fisher matrix requires no data*.

This result is extremely powerful. If the data have a (multivariate) gaussian distribution (and the errors can be correlated;  $C$  need not be diagonal), and you know how the means  $\boldsymbol{\mu}$  and the covariance matrix  $C$  depend on the parameters, you can calculate the Fisher Matrix *before you do the experiment*. The Fisher Matrix gives you the expected errors, so you know how well you can expect to do if you do a particular experiment, and you can then design an experiment to give you, for example, the best (marginal) error on the parameter you are most interested in.

Note that if the prior is not uniform, then you can simply add a ‘prior matrix’ to the Fisher matrix before inversion. Fig. 1 shows an example, where a prior from CMB experimental results has been added to a hypothetical 3D weak lensing survey.

Treat the Fisher errors as a one-way test: you might not achieve errors which are this small, but you won’t do better. So if you want to measure some quantity with an accuracy of a metre, and a Fisher analysis tells you the error bar is the size of Belgium, give up.

## Reparametrisation

Finally we mention that sometimes the gaussian approximation for the likelihood surface is not a very good approximation. With a good theoretical model, it is possible to make nonlinear transformations of the parameters to make the likelihood more gaussian. See [15] for more details.

## iCosmo: a great resource

icosmo.org has a web-based calculator for Fisher matrices for cosmology. You can download results, or even the source code, to compute Fisher matrices for various experiments in lensing, BAOs, supernovae etc, or custom-design your own survey. It also gives many other useful things, such as lensing power spectra, angular diameter distances etc.

## 4 Numerical methods

If the problem has only two or three parameters, then it may be possible to evaluate the likelihood on a sufficiently fine grid to be able to locate the peak and estimate the errors. If the dimensionality of the parameter space is very large, then, as the number of grid points grows exponentially with dimension, it becomes rapidly unfeasible to do it this way. In fact, it's very inefficient to do this anyway, as typically most of the hypervolume has very small likelihood so is of little interest. There are various ways to sample the likelihood surface more efficiently, concentrating the points more densely where the likelihood is high.

The most common method in use is Monte Carlo Markov Chain (MCMC). We will also cover here a relatively new method (to cosmology), Hamiltonian Monte Carlo, which seems more efficient in cases studied.

### 4.1 Monte Carlo Markov Chain (MCMC) method

The aim of MCMC is to generate a set of points in the parameter space whose distribution function is the same as the *target density*, in this case the likelihood, or more generally the posterior. MCMC makes random drawings, by moving in parameter space in a Markov process - i.e. the next sample depends on the present one, but not on previous ones. By design, the resulting Markov Chain of points samples the posterior, such that the density of points is proportional to the target density (at least asymptotically), so we can estimate all the usual quantities of interest from it (mean, variance, etc). The number of points required to get good estimates is said to scale linearly with the number of parameters, so very quickly becomes much faster than grids as the dimensionality increases. In cosmology, we are often dealing with around 10-20 parameters, so MCMC has been found to be a very effective tool.

The target density is approximated by a set of delta functions (you may need to normalise)

$$p(\boldsymbol{\theta}) \simeq \frac{1}{N} \sum_{i=1}^N \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_i) \quad (25)$$

from which we can estimate any integrals (such as the mean, variance etc.):

$$\langle f(\boldsymbol{\theta}) \rangle \simeq \frac{1}{N} \sum_{i=1}^N f(\boldsymbol{\theta}_i). \quad (26)$$

The basic procedure to make the chain is to generate a new point  $\boldsymbol{\theta}^*$  from the present point  $\boldsymbol{\theta}$  (by taking some sort of step), and accepting it as a new point in the chain with a probability which depends on the ratio of the new and old target densities. The distribution of steps is called the *proposal distribution*. The most popular algorithm is the *Metropolis-Hastings* algorithm, where the probability of acceptance is

$$p(\text{acceptance}) = \min \left[ 1, \frac{p(\boldsymbol{\theta}^*)q(\boldsymbol{\theta}^*|\boldsymbol{\theta})}{p(\boldsymbol{\theta})q(\boldsymbol{\theta}|\boldsymbol{\theta}^*)} \right] \quad (27)$$

where the proposal distribution function is  $q(\boldsymbol{\theta}^*|\boldsymbol{\theta})$  for a move from  $\boldsymbol{\theta}$  to  $\boldsymbol{\theta}^*$ .

If the proposal distribution is symmetric (as is often the case), the algorithm simplifies to the *Metropolis algorithm*:

$$p(\text{acceptance}) = \min \left[ 1, \frac{p(\boldsymbol{\theta}^*)}{p(\boldsymbol{\theta})} \right]. \quad (28)$$

- Choose a random initial starting point in parameter space, and compute the target density.
- Repeat:
- Generate a step in parameter space from a *proposal distribution*, generating a new trial point for the chain.
- Compute the target density at the new point, and accept it (or not) with the Metropolis-Hastings algorithm.
- If the point is not accepted, *the previous point is repeated in the chain*<sup>5</sup>.
- End Repeat:

The easy bits: this is trivial to code - you might just take a top-hat proposal distribution in each parameter direction, and it should work. The harder parts are (even in the tophat case): choosing an efficient proposal distribution; dealing with *burn-in* and *convergence*.

<sup>5</sup> It is a common mistake to neglect to do this



## Proposal distribution

If the proposal distribution is small, in the sense that the typical jump is small, then the chain may take a very long time to explore the target distribution, and it will be very inefficient. Since the target density hardly changes, almost all points are accepted, but it still takes forever. This is an example of poor *mixing*. If the proposal distribution is too large, on the other hand, then the parameter space is explored, but the trial points are often a long way from the peak, at places where the target density is low. This is very inefficient as well, since they are almost always rejected by the Metropolis-Hastings algorithm. So what is best? You might expect the chain to do well if the proposal distribution is ‘about the same size as the peak in the target density’, and you would be right. In fact, a very good option is to draw from a multivariate gaussian with the Fisher matrix as Hessian. However, having said that, if you want something quick to code (but sub-optimal), a top-hat of about the right dimensions will do a decent job. You might have to run a preliminary chain or two to get an idea of the size of the target, but the rules say you are not allowed to change the proposal distribution in a chain, so once you have decided you must throw away your chains and begin again.

## Burn-in and convergence

Theory indicates that the chain should fairly sample the target distribution once it has converged to a stationary distribution. This means that the early part of the chain (the ‘burn-in’ are ignored, and the dependence on the starting point is lost. *It is vitally important to have a convergence test.* Be warned that the points in a MCMC chain are correlated, and the chain can appear to have converged when it has not (one can reduce this problem by evaluating the correlation function of the points, and ‘thinning’ them by (say) taking only every third point (or whatever is suggested by the correlation analysis). Fig. 2 shows one such example. The classic test is the *Gelman-Rubin* (1992) convergence criterion. Start  $M$  chains, each with  $2N$  points, starting at well-separated parts of parameter space. In this example, the first  $N$  are discarded as burn-in.

The idea is that you have two ways to estimate the mean of the parameters - either treat the combined chains as a single dataset, or look at the means of each chain. If the chains have converged, these should agree within some tolerance.

Following [26], let  $\theta_I^J$  represent the point in parameter space in position  $I$  of chain  $J$ . Compute the mean of each chain ( $J$ ):

$$\bar{\theta}^J \equiv \frac{1}{N} \sum_{I=1}^N \theta_I^J \quad (29)$$

and the mean of all the chains

$$\bar{\theta} \equiv \frac{1}{NM} \sum_{I=1}^N \sum_{J=1}^M \theta_I^J. \quad (30)$$

The chain-to-chain variance  $B$  is

$$B = \frac{1}{(M-1)} \sum_{J=1}^M (\bar{\theta}^J - \bar{\theta})^2 \quad (31)$$

and the average variance of each chain is

$$W = \frac{1}{M(N-1)} \sum_{I=1}^N \sum_{J=1}^M (\theta_I^J - \bar{\theta}^J)^2. \quad (32)$$

Under convergence,  $W$  and  $B/N$  should agree.

The weighted estimate of the variance,

$$\sigma^2 = \frac{N-1}{N} W + \frac{B}{N} \quad (33)$$

overestimates the true variance if the starting distribution is overdispersed.


Accounting for the variance of the means gives an estimator of the variance

$$V = \sigma^2 + \frac{B}{NM} \quad (34)$$

$V$  is an overestimate, and  $W$  is an underestimate. The ratio of the two estimates is

$$\hat{R} = \frac{\left(\frac{N-1}{N}\right) + B\left(1 + \frac{1}{M}\right)}{W}. \quad (35)$$

$R$  should approach unity as convergence is achieved. How close to  $R = 1$ ? Opinions differ; I have seen suggestions to run the chain until the values of  $\hat{R}$  are always  $< 1.03$ , or  $< 1.2$ , but a proof would be nice.



**Fig. 2.** Examples of unconverged chains. The left panel suggests the chain has converged, but the right panel shows the chain and a second one, also apparently converged, but showing clearly different distributions. From Verde et al. ApJS, 148, 195 (2003).

## CosmoMC

Excellent resource. MCMC sampler with cosmological data (CMB + support for LSS, SNe).  
<http://cosmologist.info/cosmomc/>

For more details on MCMC in general, see [6, 17, 27].

## 4.2 Hamiltonian Monte Carlo

As we've seen, the proposal distribution has to be reasonably finely tuned to ensure good mixing, but not be too inefficient. What we would really like is to be able to take rather big jumps, so the chain is well-mixed, but in such a way that the probability of acceptance of each point is still high. Hamiltonian (or Hybrid) Monte Carlo (HMC) tries to do this, via a rather clever trick. In practice, it seems that it can be about  $M$  times faster than MCMC in  $M$  dimensions. Typical applications report that 4 times shorter chains give the same accuracy as MCMC. Originally developed for particle physics [3], there is a nice exposition in astrophysics by Hajian [7].

HMC works by sampling from a *larger* parameter space than we want to explore, by introducing  $M$  *auxiliary variables*, one for each parameter in the model. To see how this works, imagine each of the parameters in the problem as a coordinate. HMC regards the target distribution which we seek as an effective potential in this coordinate system, and for each coordinate it generates a generalised momentum. i.e. it expands the parameter space from its original 'coordinate space' to a phase space, in which there is a so-called 'extended' target distribution. It tries to sample from the extended target distribution in the  $2M$  dimensions. It explores this space by treating the problem as a dynamical system, and evolving the phase space coordinates by solving the dynamical equations. Finally, it ignores the momenta (marginalising, as in MCMC), and this gives a sample of the original target distribution.

There are several advantages to this approach. One is that the probability of acceptance of a new point is high - close to unity; secondly, the system can make big jumps, so the mixing is better and the convergence faster. In doing the big jumps, it does some additional calculations, but these do not involve computing the likelihood, which is typically computationally expensive.

Let us see how it works. If the target density in  $M$  dimensions is  $p(\boldsymbol{\theta})$ , then we define a potential

$$U(\boldsymbol{\theta}) \equiv -\ln p(\boldsymbol{\theta}). \quad (36)$$

For each coordinate  $\theta_\alpha$ , we generate a momentum  $u_\alpha$ , conveniently from a normal distribution with zero mean and unit variance, so the  $M$ -dimensional momentum distribution is a simple multivariate gaussian which we denote  $\mathcal{N}(\mathbf{u})$ . We define the kinetic energy

$$K(\mathbf{u}) \equiv \frac{1}{2} \mathbf{u}^T \mathbf{u}, \quad (37)$$

and the Hamiltonian is

$$H(\boldsymbol{\theta}, \mathbf{u}) \equiv U(\boldsymbol{\theta}) + K(\mathbf{u}). \quad (38)$$

The trick is that we generate chains to sample the *extended target density*

$$p(\boldsymbol{\theta}, \mathbf{u}) = \exp[-H(\boldsymbol{\theta}, \mathbf{u})]. \quad (39)$$

Since this is separable,

$$p(\boldsymbol{\theta}, \mathbf{u}) = \exp[-U(\boldsymbol{\theta})] \exp[-K(\mathbf{u})] \propto p(\boldsymbol{\theta}) \mathcal{N}(\mathbf{u}) \quad (40)$$

and if we then marginalise over  $\mathbf{u}$  by simply ignoring the  $\mathbf{u}$  coordinates attached to each point in the chain (just as in MCMC), the resulting marginal distribution samples the desired target distribution  $p(\boldsymbol{\theta})$ . This is really very neat.

The key is that if we *exactly* solve the Hamiltonian equations

$$\begin{aligned} \dot{\theta}_\alpha &= u_\alpha \\ \dot{u}_\alpha &= -\frac{\partial H}{\partial \theta_\alpha} \end{aligned} \quad (41)$$

then  $H$  remains invariant, so the extended target density is always the same, and the acceptance is unity. Furthermore, we can integrate the equations for a long time if we wish, decorrelating the points in the chain.

There are several issues to consider.

- We seem to need the target density to define the potential, but this is what we are looking for. We need to approximate it.
- The aim is to do this fast, so we do not want to do many operations before generating a new sample. An easy way to achieve this is to employ a simple integrator (e.g. leap-frog; even Euler's method might do) and take several fairly big steps before generating a new point in the chain.
- We need to ensure we explore the extended space carefully.

The result of the two approximations is that  $H$  will not be quite constant. We deal with this as in MCMC by using the Metropolis algorithm, accepting the new point  $(\boldsymbol{\theta}^*, \mathbf{u}^*)$  with a probability

$$\min\{1, \exp[-H(\boldsymbol{\theta}^*, \mathbf{u}^*) + H(\boldsymbol{\theta}, \mathbf{u})]\}, \quad (42)$$


otherwise we repeat the old point  $(\boldsymbol{\theta}, \mathbf{u})$  as usual.

How do we approximate? We want the gradients to be cheap to compute. It is usual to generate an approximate analytical potential by running a relatively short MCMC chain, computing the covariance of the points in the chain, and approximating the distribution by a multivariate gaussian by a multivariate gaussian with the same covariance. The gradients are then consequently easy to compute analytically.

The last point is that if we change the momentum only with Hamilton's equations of motion, we will restrict ourselves to a locus in phase space, and the target distribution will not be properly explored. To avoid this, a new momentum is generated randomly when each point in the chain is generated. The art is to choose a good step in the integration, and the number of steps to take before generating a new point. Perhaps unsurprisingly, choosing these such that the new point differs from the previous one by about the size of the target peak works well. Thinning can be performed, and a convergence test must still be applied. Fig. 3 shows a comparison between HMC and MCMC for a simple case of a 6D gaussian target distribution, from [7].

## 5 Model Selection

Model selection is in a sense a higher-level question than parameter estimation. In parameter estimation, one assumes a theoretical model within which one interprets the data, whereas in model selection, one wants to



**Fig. 3.** A comparison of HMC sampling (top) and MCMC sampling (bottom). Note that the computer time required to generate each point in the HMC sampling will be larger than that of MCMC, so the actual gains are less than appears. From [7].

know which theoretical framework is preferred, given the data (regardless of the parameter values). The models may be completely different (e.g. compare Big Bang with Steady State, to use an old example), or variants of the same idea. E.g. comparing a simple cosmological model where the Universe is assumed flat and the perturbations are strictly scale-invariant ( $n = 1$ ), with a more general model where curvature is allowed to vary and the spectrum is allowed to deviate from scale-invariance. The sort of question asked here is essentially ‘Do the data require a more complex model?’. Clearly in the latter type of comparison  $\chi^2$  itself will be of no use - it will always reduce if we allow more freedom. There are frequentist ways to try and answer these questions, but we are all by now confirmed Bayesians<sup>6</sup>, so will approach it this way.

### 5.1 Bayesian evidence

The Bayesian method to select between models (e.g. General Relativity, & modified gravity) is to consider the Bayesian evidence ratio. Essentially we want to know if, given the data, there is evidence that we need to expand the space of gravity models beyond GR. Assuming non-committal priors for the models (i.e. the same a priori probability), the probability of the models given the data is simply proportional to the evidence.

We denote two competing models by  $M$  and  $M'$ . We assume that  $M'$  is a simpler model, which has fewer ( $n' < n$ ) parameters in it. We further assume that it is *nested* in Model  $M$ , i.e. the  $n'$  parameters of model  $M'$  are common to  $M$ , which has  $p \equiv n - n'$  extra parameters in it. These parameters are fixed to fiducial values in  $M'$ .

We denote by  $\mathbf{x}$  the data vector, and by  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}'$  the parameter vectors (of length  $n$  and  $n'$ ).

Apply Rule 1: Write down what you want to know. Here it is  $p(M|\mathbf{x})$  - the probability of the model, given the data.

The posterior probability of each model comes from Bayes’ theorem:

$$p(M|\mathbf{x}) = \frac{p(\mathbf{x}|M)p(M)}{p(\mathbf{x})} \quad (43)$$

and similarly for  $M'$ . By marginalisation  $p(\mathbf{x}|M)$ , known as the *Evidence*, is

$$p(\mathbf{x}|M) = \int d\boldsymbol{\theta} p(\mathbf{x}|\boldsymbol{\theta}M)p(\boldsymbol{\theta}|M), \quad (44)$$

which should be interpreted as a multidimensional integration. Hence the posterior relative probabilities of the two models, regardless of what their parameters are<sup>7</sup>, is

$$\frac{p(M'|\mathbf{x})}{p(M|\mathbf{x})} = \frac{p(M')}{p(M)} \frac{\int d\boldsymbol{\theta}' p(\mathbf{x}|\boldsymbol{\theta}'M')p(\boldsymbol{\theta}'|M')}{\int d\boldsymbol{\theta} p(\mathbf{x}|\boldsymbol{\theta}M)p(\boldsymbol{\theta}|M)}. \quad (45)$$

With non-committal priors on the models,  $p(M') = p(M)$ , this ratio simplifies to the ratio of evidences, called the *Bayes Factor*,

<sup>6</sup> If not, please leave the room

<sup>7</sup> If a model has no parameters, then the integral is simply replaced by  $p(\mathbf{x}|M)$

$$B \equiv \frac{\int d\boldsymbol{\theta}' p(\mathbf{x}|\boldsymbol{\theta}'M')p(\boldsymbol{\theta}'|M')}{\int d\boldsymbol{\theta} p(\mathbf{x}|\boldsymbol{\theta}M)p(\boldsymbol{\theta}|M)}. \quad (46)$$

Note that the a complicated model  $M$  will (if  $M'$  is nested) inevitably lead to a higher likelihood (or at least as high), but the evidence will favour the simpler model if the fit is nearly as good, through the smaller prior volume.

We assume uniform (and hence separable) priors in each parameter, over ranges  $\Delta\boldsymbol{\theta}$  (or  $\Delta\boldsymbol{\theta}'$ ). Hence  $p(\boldsymbol{\theta}|M) = (\Delta\boldsymbol{\theta}_1 \dots \Delta\boldsymbol{\theta}_n)^{-1}$  and

$$B = \frac{\int d\boldsymbol{\theta}' p(\mathbf{x}|\boldsymbol{\theta}', M') \Delta\boldsymbol{\theta}_1 \dots \Delta\boldsymbol{\theta}_n}{\int d\boldsymbol{\theta} p(\mathbf{x}|\boldsymbol{\theta}, M) \Delta\boldsymbol{\theta}'_1 \dots \Delta\boldsymbol{\theta}'_{n'}}. \quad (47)$$

Note that if the prior ranges are not large enough to contain essentially all the likelihood, then the position of the boundaries would influence the Bayes factor. In what follows, we will assume the prior range is large enough to encompass all the likelihood.

In the nested case, the ratio of prior hypervolumes simplifies to

$$\frac{\Delta\boldsymbol{\theta}_1 \dots \Delta\boldsymbol{\theta}_n}{\Delta\boldsymbol{\theta}'_1 \dots \Delta\boldsymbol{\theta}'_{n'}} = \Delta\boldsymbol{\theta}_{n'+1} \dots \Delta\boldsymbol{\theta}_{n'+p}, \quad (48)$$

where  $p \equiv n - n'$  is the number of extra parameters in the more complicated model.

Here we see the problem. The evidence requires a multidimensional integration over the likelihood and prior, and this may be *very* expensive to compute. There are various ways to simplify this. One is analytic - follow the Fisher approach and assume the likelihood is a multivariate gaussian, others are numerical, such as nested sampling, where one tries to sample the likelihood in an efficient way. There are others, but we will focus on these. Note that shortcuts with names such as AIC and BIC may be unreliable as they are based on the best-fit  $\chi^2$ , and from a Bayesian perspective we want to know how much parameter space would give the data with high probability. See [18] for more discussion.

## 5.2 Laplace approximation

The Bayes factor in equation (47) still depends on the specific dataset  $\mathbf{x}$ . For future experiments, we do not yet have the data, so we compute the expectation value of the Bayes factor, given the statistical properties of  $\mathbf{x}$ . The expectation is computed over the distribution of  $\mathbf{x}$  for the correct model (assumed here to be  $M$ ). To do this, we make two further approximations: first we note that  $B$  is a ratio, and we approximate  $\langle B \rangle$  by the ratio of the expected values, rather than the expectation value of the ratio. This should be a good approximation if the evidences are sharply peaked.

We also make the Laplace approximation, that the expected likelihoods are given by multivariate Gaussians. For the correct model  $M$ , the peak of the expected likelihood is located at the true parameters  $\boldsymbol{\theta}_0$ . Note, however, that for the incorrect model  $M'$ , the peak of the expected likelihood is not in general at the true parameters (see Fig. 4 for an illustration of this). This arises because the likelihood in the numerator of equation (47) is the probability of the dataset  $\mathbf{x}$  given incorrect model assumptions.


The Laplace approximation is routinely used in forecasting marginal errors in parameters, using the Fisher matrix. Clearly the approximation may break down in some cases, but for Planck, the Fisher matrix errors are reasonably close to (within 30% of) those computed with Monte Carlo Markov Chains. If we assume that the posterior probability densities are small at the boundaries of the prior volume, then we can extend the integrations to infinity, and the integration over the multivariate Gaussians can be easily done.

The final expression (see [10] for a derivation) for the expected Bayes factor is then

$$\langle B \rangle = (2\pi)^{-p/2} \frac{\sqrt{\det \mathbf{F}}}{\sqrt{\det \mathbf{F}'}} \exp\left(-\frac{1}{2} \delta\boldsymbol{\theta}_\alpha \mathbf{F}_{\alpha\beta} \delta\boldsymbol{\theta}_\beta\right) \prod_{q=1}^p \Delta\boldsymbol{\theta}_{n'+q}. \quad (49)$$

where the offsets are given by  $\delta\boldsymbol{\theta}_\alpha = \delta\boldsymbol{\theta}'_\alpha$  for  $\alpha \leq n'$  and  $\delta\boldsymbol{\theta}_\alpha = \delta\psi_{\alpha-n'}$  for  $\alpha > n'$ . Note that  $\mathbf{F}$  and  $\mathbf{F}^{-1}$  are  $n \times n$  matrices,  $\mathbf{F}'$  is  $n' \times n'$ , and  $\mathbf{G}$  is an  $n' \times p$  block of the full  $n \times n$  Fisher matrix  $\mathbf{F}$ . The expression we find is a specific example of the Savage-Dickey ratio (see e.g. [25]). For a nested model with a single additional parameter  $\boldsymbol{\theta}_i$ ,  $\langle B \rangle = p(\boldsymbol{\theta}_i|\mathbf{x})/p(\boldsymbol{\theta}_i)$ .

Here we explicitly use the Laplace approximation to compute the offsets in the parameter estimates which accompany the wrong choice of model, and compute the evidence ratio explicitly. Finally, note that this is the expected evidence ratio (nearly); it does not address the issue of what the distribution of evidence ratios should be.



**Fig. 4.** Illustrating how assumption of a wrong parameter value can influence the best-fitting value of other model parameters. Ellipses represent iso-likelihood surfaces, and here in the simpler model, the parameter on the horizontal axis is assumed to take the value given by the vertical line. Filled circles show the true parameters in the more complicated model, and the best-fit parameters in the simpler model. From [9].

Note that the ‘Occam’s razor’ term, common in evidence calculations, is to some extent encapsulated in the term  $(2\pi)^{-p/2} \frac{\sqrt{\det \mathbf{F}}}{\sqrt{\det \mathbf{F}'}}$ , multiplied by the prior product: models with more parameters are penalised in favour of simpler models, unless the data demand otherwise. In cases where the Laplace approximation is not a good one, other techniques must be used, at more computational expense.

It is perhaps worth remarking that Occam’s razor appears not to be fully incorporated into this term, as can be seen by considering a situation where the data do not depend at all on an additional parameter. In this case, the Bayesian evidence ratio is unity, so disappointingly no preference is shown at all.

As an example, consider testing General Relativity against other gravity theories, which predict a different growth rate of perturbations,  $d \ln \delta / d \ln \Omega_m = \gamma$ , where  $\gamma = 0.55$  for GR, and (for example)  $\gamma = 0.68$  for a flat DGP braneworld model. This can be probed with weak lensing, and we ask the question do the data favour a model where  $\gamma$  is a free parameter, rather than being fixed at 0.55?


We take a prior range  $\Delta\gamma = 1$ , and we ask the question of how different the growth rate of a modified-gravity model would have to be for these experiments to be expected to favour a relaxation of the gravity model from General Relativity. This is shown in Fig.5. It shows how the expected evidence ratio changes with progressively greater differences from the General Relativistic growth rate. We see that a next-generation weak lensing survey could even distinguish ‘strongly’  $\delta\gamma = 0.048$ . Note that changing the prior range  $\Delta\gamma$  by a factor 10 changes the numbers by  $\sim 0.012$ , so the dependence on the prior range is rather small.

If one prefers to ask a frequentist question, then a combination of WL+*Planck*+BAO+SN should be able to distinguish  $\delta\gamma = 0.13$ , at  $10.6\sigma$ . Alternatively, one can calculate the expected error on  $\gamma$  [1] within the extended model  $M$ . In this section, we are asking a slightly different question of whether the data demand that a wider class of models needs to be considered at all, rather than estimating a parameter within that wider class.

The case for a large, space-based 3D weak lensing survey is strengthened, as it offers the possibility of conclusively distinguishing Dark Energy from at least some modified gravity models.

### 5.3 Numerical sampling methods


In order to compute the evidence numerically, the prior volume needs to be sampled, in much the same way as in parameter estimation, except the requirements are slightly more stringent. MCMC could be used, although there are claims that it is not good at exploring the parameter space adequately. I find these claims puzzling, as in evidence calculations we are doing an  $N$ -dimensional integral, which is not so different from doing the  $(N - 1)$ -dimensional integral in MCMC to get marginal errors. However, here are a couple of other sampling techniques.



**Fig. 5.** The expected value of  $|\ln(B)|$  from a future large-scale deep weak lensing survey, as might be done with Euclid, JDEM or LSST, in combination with CMB constraints from *Planck*, as a function of the difference in the growth rate between the modified-gravity model and General Relativity. The crossover at small  $\delta\gamma$  occurs because Occam's razor will favour the simpler (General Relativity) model unless the data demand otherwise. To the left of the cusp, GR would be likely to be preferred by the data. The dotted vertical line shows the offset of the growth factor for the flat DGP model. The descriptors are the terminology of Jeffreys (1961) [12]. From [9]

### Nested sampling

Other sampling methods can also be very effective (see e.g. [25, 2, 11, 19]). Nested sampling was introduced by Skilling [23]. One samples from the prior volume, and gradually concentrates more points near the peak of the likelihood distribution, by repeatedly replacing the point with the lowest target density by one drawn strictly from the prior volume with higher target density. This has proved effective for cosmological model selection. I will not go into details here, except to say that the key is in drawing the new point from a suitable subset of the prior volume. This must be increasingly smaller as the points get more confined, otherwise the trial points will be rejected with increasing frequency and it becomes very inefficient, but it must also be a large enough subset that the entire prior volume above the lowest target density is almost certainly included inside. For further details, see the original papers. Note that, for multimodal target distributions, a modification called MultiNest exists [4]. CosmoNest and MultiNest are publicly-available additions to CosmoMC.



**Fig. 6.** The Nested Sampling algorithm. From [19]

## 6 Summary

We have explored the use of (principally) Bayesian methods for parameter estimation and model selection, and looked at some common numerical methods for performing these tasks. We have also shown how it is possible to work out the (minimum) errors on parameters, and on model probabilities, in advance of doing an experiment, to see if it is worthwhile. These are the common Fisher matrix approach for parameter estimation, and the expected evidence for model selection, which requires no more than the Fisher matrix to compute.

## References

1. Amendola L., Kunz M., Sapone D., JCAP, 04, 013 (2008)
2. Beltran M., Garcia-Bellido J., Lesgourgues J., Liddle A. R., Slosar A., 2005, Phys. Rev., D71, 063532
3. Duan et al., Phys. Lett. B195, 216 (1987)
4. Feroz F., Hobson M., Bridges M., astro-ph/0809.3437 (2008)
5. Fisher R. A., J. Roy. Stat. Soc., 98, 39 (1935)
6. Gilks W.R., Richardson S., Spiegelhalter D.J., *Markov chain Monte Carlo in practice*, Chapman & Hall, London (1996)
7. Hajian A., astro-ph/0608679 (2006)
8. Hamilton A.J.S., astro-ph/0503603 (2005)
9. Heavens A.F., Kitching T., Verde L., MNRAS, 380, 1029 (2007).
10. Heavens A.F., arXiv:0906.0664 (2009)
11. Hobson M.P., Bridle S.L., Lahav O., 2002, MNRAS, 335, 377
12. Jeffreys H., 1961, *Theory of Probability*, Oxford University Press (Oxford, UK)
13. Kendall M. G., Stuart A., *The Advanced Theory of Statistics, Volume II*, Griffin, London (1969)
14. Kenney, J. F. & Keeping, E. S., *Mathematics of Statistics, Part II*, 2nd ed. (Van Nostrand, New York) (1951)
15. Kosowsky A., Milosavljevic M., Jimenez R., Phys. Rev. D66, 3007 (2002).
16. Lazarides, G., Ruiz de Austri, R., Trotta, R., Phys. Rev. D70, 123527
17. Lewis A., Bridle S., PRD, 66, 103511 (2002)
18. Liddle A., MNRAS, 377, 74 (2007).
19. Mukherjee P., Parkinson D., Liddle A.R., 2006, ApJ, 638, L51
20. Press W., et al, *Numerical Recipes in Fortran*, CUP, Cambridge, U.K. (1992)
21. Serra P., Heavens A., Melchiorri A., MNRAS 379, 169 (2007)
22. Sivia D.S., Skilling J., *Data Analysis; A Bayesian tutorial*, OUP, Oxford (2006)
23. Skilling J., 2004, available at  
<http://www.inference.phy.cam.ac.uk/bayesys>
24. Tegmark M., Taylor, A.N., Heavens A.F., ApJ, 480, 22 (1997) (TTH)
25. Trotta R., 2007, astro-ph/0703063
26. Verde L., astro-ph/0712.3028 (2007)
27. Verde L., et al, ApJS, 148, 195 (2003)