CHAPTER 14

Bayesian approaches to the quantitative genetic analysis of natural populations

Michael B. Morrissey, Pierre de Villemereuil, Blandine Doligez and Olivier Gimenez

14.1 Introduction

Evolutionary quantitative genetic analysis of natural populations is proving to be highly rewarding, but also comes with enormous challenges. Parameters that have always been regarded as difficult to estimate in the laboratory, for example genetic correlations, are even more difficult in data that contains the 'real' noise of nature. It is therefore very important to consider the best models that we can use for the study of data from natural populations, but also to consider the statistical uncertainty inherent to the estimates yielded by these models. Natural populations also present the quantitative geneticist with additional complications; in particular, it may become increasingly important to explicitly consider the process of observing data in conjunction with inferences about underlying biological processes. For example, evaluating life histories when individuals are not perfectly observable will help to be cautious about genetic inferences. Bayesian techniques offer the empiricist some of the most promising ways of dealing with complicated and noisy data.

In this first section we take a very close look at estimation of heritability in a simple breeding experiment. This is not an analysis of a natural population, but provides a simple example that illustrates a range of non-trivial details with which any empiricist must familiarise him- or herself before putting Bayesian methods into practice. In short, the idea here is to turn the usual mode of presentation of a statistical method on its head. Rather than starting from a completely trivial model and building toward complex and scientifically interesting analyses, we are starting from an assumption that the reader is sufficiently familiar with the basic biological principle of inferring the genetic basis of traits from similarity among relatives (see Chapter 2, Postma). Given this biologist's view of the flow of phenotype data through a model to make genetic inferences based on similarity of relatives, we hope that the important statistical and Bayesian aspects of interpreting the models will be as intuitive as possible, and we will deal with these aspects as they arise. In doing so we hope to overcome the biggest obstacle to realising the potential benefits of Bayesian quantitative genetic analysis in the wild: getting off the ground.

Our biologist's view of a simple quantitative genetic analysis, and the subsequent more developed applications to natural populations in the next section, do not provide any comprehensive guide to either Bayesian philosophies or methodologies. General practical texts include *Bayesian data analysis* (Gelman *et al.* 2004), *Data*

Quantitative Genetics in the Wild. Edited by Anne Charmantier, Dany Garant, and Loeske E. B. Kruuk

analysis using regression and multilevel/hierarchical models (Gelman & Hill 2007), and *Doing bayesian data analysis* (Kruschke 2011). Those with an ecological background who are interested in incorporating Bayesian, 'modular' or hierarchical analyses, and potentially quantitative genetic methods into their work may find *Hierarchical modeling and inference in ecology* (Royle & Dorazio 2008) to be a useful introduction to Bayesian Markov Chain Monte Carlo (MCMC)-based analysis. An important general resource for MCMC analysis is the book *Handbook of Markov Chain Monte Carlo* (Brooks *et al.* 2011). A vastly more complete introduction to Bayesian quantitative genetic analysis, which includes many ideas that would greatly benefit the analysis of natural populations, is *Likelihood, Bayesian, and MCMC methods in quantitative genetics* (Sorensen & Gianola 2002). O'Hara *et al.* (2008) 'Bayesian approaches in evolutionary quantitative genetics' is a useful review that is less specifically focused on evolutionary problems in natural populations, but touches on similar and complimentary themes to ours.

In the second section, we explore some specific cases where extension of the approaches and thinking from the first section is currently allowing informative, cutting-edge quantitative genetic analyses of natural populations. In each example, the inherent flexibility of Bayesian approaches, and available Bayesian tools, is particularly important in allowing more direct inference of key evolutionary parameters than is often possible in frequentist frameworks. The goal of the current section is to suggest several ways in which Bayesian analyses can potentially provide insight into current microevolutionary problems. This section is intended to be less didactic in terms of details of implementation. We describe several types of analyses that we expect to become increasingly common in the near future. Whilst we seek to explain the established aspects of their implementation, we do not intend that our treatment should in any way be regarded as a guide to 'best practices', because these are developing at a great pace. Rather, as before, we seek primarily to outline the utility and flexibility that Bayesian analysis can bring to the quantitative genetic analysis of natural populations. Each of the examples we discuss illustrates a specific way in which the Bayesian toolkit allows specific inference of evolutionary parameters that would be very difficult (but probably in no case impossible) to obtain otherwise.

14.2 Putting Bayesian methods into practice: a guided tour of a simple example

14.2.1 Heritability of morphological traits in crickets

We will be very explicit about i) how the maths in the example relate to quantitative genetic parameters and ii) how the specific maths can be implemented, using the BUGS programming language (Lunn *et al.* 2000), implemented with the software JAGS (Plummer 2003). This depth will subsequently prove valuable in the next section where we refit the model in several different ways to get a feel for some important details about prior specification, when we introduce the animal model in a Bayesian implementation, and more generally, as we move through the more interesting examples throughout the second section of the chapter.

We analyse phenotypic data from a quantitative genetic experiment on field crickets, *Teleogryllus oceanicus*, by Simmons and Garcia-Gonzalez (2007) who mated 30 males to a total of 84 females, and measured 378 female offspring for a range of traits, including pronotum length and ovary mass. The goal here is to estimate additive genetic variances, heritabilities, and to control for and characterise any potentially confounding sources of variances, such as that arising from maternal effects. The main trick is to characterise the amount of variation due to sires. Four times the sire variance is the additive genetic variance, and the quotient of the additive genetic variance and the phenotypic variance is the heritability. The simplest mixed model with which to analyze this 'dams within sires' breeding experiment is

$$y_i = \mu + s_i + d_i + e_i \qquad (14.1)$$

where y_i is the phenotype of individual i, μ is the population mean, s_i and d_i represent the effects of the dam and sire of individual i, and e are residual errors. Because we are interested in the variance among s and d values, not necessarily

the effects of each dam and sire in isolation, we model them as random effects. What this means is that we assume that s_i, d_i, and similarly e_i values come from normal distributions, the variances of which are the parameters of interest, and which are estimated. This can be written $s_i \sim N(0, \sigma_s^2)$, $d_i \sim N(0, \sigma_d^2)$, and $e_i \sim N(0, \sigma_e^2)$, where σ_x^2 represents the variance among effects x. More fully, we could write the full likelihood as

$$\begin{aligned}\ell_i(\mu, s_i, d_i, \sigma_\epsilon^2) &= p(\mathbf{y}|\mu, s_i, d_i, \sigma_e^2) \\ &= \prod_{i=1}^{N_i} N(y_i|\mu + s_i + d_i, \sigma_e^2)\end{aligned} \tag{14.2a}$$

$$\ell(\sigma_d^2) = p(\mathbf{d}|\sigma_d^2) = \prod_{j=1}^{N_d} N(d_j|0, \sigma_d^2) \tag{14.2b}$$

$$\ell(\sigma_s^2) = p(\mathbf{s}|\sigma_s^2) = \prod_{k=1}^{N_s} N(s_k|0, \sigma_s^2) \tag{14.2c}$$

$$\ell(\mu, \sigma_{e(y)}^2, \sigma_d^2, \sigma_s^2) = p(\mathbf{y}|\mu, s_i, d_i, \sigma_e^2) \times p(\mathbf{d}|\sigma_d^2) \times p(\mathbf{s}|\sigma_s^2) \tag{14.2d}$$

where ℓ represents the likelihoods of parameters, p represents the probability of given observed data (i.e. y_i) or effects of dams and sires, i.e. the vectors $\mathbf{d}$ and $\mathbf{s}$. The left-hand side of each expression in Eq. 14.2 represents the likelihood of parameter estimates, which is equated to the probability of some data, i.e. y in Eq. 14.2a or the unobservable sire and dam effects in Eqs. 14.2b,c, or all jointly in 14.2d. The right-most expressions in Eq. 14.2a,b,c are the core of the model. Here the probabilities are specified in terms of products of normal density functions associated with each data observation or random effect.

Eq. 14.2 may seem like an unnecessarily complex way of representing the mixed model that we managed with a single line in Eq. 14.1. However, this representation lies at the core of the problem, either in a frequentist likelihood, or in a Bayesian analysis. In order to analyse genetic parameters in the cricket dataset in a Bayesian framework, we have to do two things. First, we have to come up with priors for the parameters, and second, we have to think about implementation. To start with, we will try to apply the simplest priors we can think of. The parameter μ can in principle take any real value, so a very wide (high variance) normal density is simple and thorough. For the variance of each effect, i.e. the sire, dam, and residual sources of variation, any positive value is permissible. We will start by allowing all values greater than zero, and up to some large value (ideally, this should not be informed by the data in any way; in this case, making sure that the upper limit is substantially larger than the observed variance should be pragmatic) to be equally likely, i.e. we will apply uniform prior densities on the standard deviation (SD). Formally, we could express these prior choices as

$$\mu \sim N(0, 1000) \tag{14.3a}$$

$$\sigma_s \sim U(0, \sigma_{big}) \tag{14.3b}$$

$$\sigma_d \sim U(0, \sigma_{big}) \tag{14.3c}$$

$$\sigma_e \sim U(0, \sigma_{big}) \tag{14.3d}$$

where $N()$ is defined as above, and $U(x, y)$ represents a uniform density with minimum x and maximum y; σ_{big} is thus an arbitrarily large upper limit.

In BUGS code, the model is thus:

```
model{
  #priors
  mu~dnorm(0,0.001)  # bugs works with precision, i.e., 1/variance
  sigma_s~dunif(0,sigma_big)
  sigma_d~dunif(0,sigma_big)
  sigma_e~dunif(0,sigma_big)

  #random effects
  for(j in 1:N_d){
     d[j]~dnorm(0,1/sigma_d^2)
   }
   for(k in 1:N_s){
```

```
13      s[k]~dnorm(0,1/sigma_s^2)
14    }
15
16    #data
17    for(i in 1:N_i){
18      y[i]~dnorm(mu+s[sire[i]]+d[dam[i]],1/sigma_e^2)
19    }
20  }
```

Computer code can be as intimidating as math. However, having already taken a moment to write out the model in full (i.e. in Eq. 14.2), and to formally define some priors on the free parameters (i.e. in Eq. 14.3), all we have here is a statement of the same model in a different syntax. Lines 3 to 6 state that the priors have the normal and uniform distributions we stated. Lines 10 and 13 correspond to the expressions in Eq. 14.2b,c, and line 18 similarly correspondso Eq. 14.2a. The information about the BUGS language thats needed to understand fully this code is that `x~a(b,c)` means 'x is sampled from distribution `a` with parameters `b` and `c`'. In the case of the uniform distribution (`dunif`), the parameters are the upper and lower bounds of the distribution, and for the normal distribution, the parameters are the mean and the precision, which is the inverse of the variance. Looping over each datum and random effect level is accomplished with 'for' loops, where `for(a in b:c){d}` means 'sequentially assign `a` to all inclusive integer values between `b` and `c`, and given these values of `a`, do `d`'. Indexing of elements of vectors is accomplished with square brackets. In the code above, several constants have to be provided to allow it to run. These are the numbers of dams, sires, and phenotypic observations (`N_d`, `N_s`, and `N_i`; note that indexes remain consistent with Eq. 14.2), a vector of the phenotypic observations (`y`) and vectors indicating which dam and sire is associated with each observation (`dam` and `sire`). Box 14.1 provides a brief overview of how models can be put to work, once coded in this way, using MCMC algorithms.

Box 14.1 Bayesianism

In statistical inference, Bayesianism is a paradigm in which (and contrary to frequentism), one consider parameters (hereafter θ) as random variables. In this sense, Bayesians are not interested in infering a point estimate (based on maximum likelihood for example) from the data Y, but a *posterior distribution* $P(\theta|Y)$. In order to do that, they use Bayes' inferential theorem:

$$P(\theta|Y) = \frac{P(Y|\theta)\,P(\theta)}{P(Y)}$$

where $P(Y|\theta)$ is the *likelihood* of the data given the model and the parameters, $P(\theta)$ is the *prior distribution* on the parameters and $P(Y)$ is a scaling constant. In essence, Bayesianism considers any inference as an *update* from your *prior belief* of what the parameters are to your *posterior belief* of what values are more probable now that you have analysed the data.

Different flavours of the MCMC algorithm represent the most popular Bayesian estimation algorithm, because they allow posterior distributions to be sampled for very arbitrary models. Thus, when models are becoming too complex for maximum likelihood estimation, Bayesianism and MCMC sampling are often used as alternative resources. MCMC algorithms yield samples of the parameters of a model in proportion to their posterior probability. For example, a well-conducted Bayesian analysis of a dataset wherein $V_A = 2$ would ideally yield many outputs of V_A in the vicinity of 2, and relatively fewer outputs of values substantially lower or greater than 2. The amount of evidence available will determine just how much those 'values substantially lower or greater than 2' actually differ from 2. Such samples are used to calculate posterior properties of the parameter (mean, median, variance, credible interval . . .), as well as to obtain posterior distributions of derived parameters of interest, such as h^2.

14.2.2 Posterior transformation to make inference of parameters of direct biological interest

Our main question was: what are the heritabilities of pronotum length and ovary mass? Given the posterior distributions of the sire models of the traits, we need only to apply the standard relationship between the sire variance and the additive genetic variance, and between the additive genetic and phenotypic variances, to obtain inference of the heritability. The most probable values of the genetic variance and heritability will be the modal values of their posterior distributions, and the posterior distributions in full provide inference of the probabilities, given the data, the model structure, and the priors, that the true values of the genetic variance and heritability take any other values. The posterior distribution of the additive genetic variance, obtained by applying $\sigma^2_{\mathrm{A}} = 4\sigma^2_s$ to each sample of the posterior distribution is shown in Figure 14.1a,c, and the posterior distribution heritability, based on $\sigma^2_{\mathrm{P}} = \sigma^2_s + \sigma^2_d + \sigma^2_e$ is given in Figure 14.1b,d.

Thus, inference of genetic parameters can be very simple, given a fitted Bayesian model (Table 14.1). Parameter estimates are similarly easily obtained in a frequentist framework, but inference of the statistical support for estimates in a frequentist framework is much harder to obtain. For example, the implementation of a sire model with the software ASReml-R (Gilmour *et al.* 1999) is

```
asreml(fixed = pro ~ 1, random = ~sire + dam, data = crickets)
```

which is much easier than the implementation route we took above. Note however that Bayesian implementation with MCMCglmm (Hadfield 2010), an R-package for fitting Bayesian generalised linear mixed models by MCMC, would be similarly simple. However, we would see and understand less of what was happening under the hood, and some of the customising of the sire model in the next section would not be possible.

The restricted maximum likelihood (REML) solution of the model for pronotum length gives the parameter estimates $\sigma^2_s = 0.027$, $\sigma^2_d = 0.012$, and

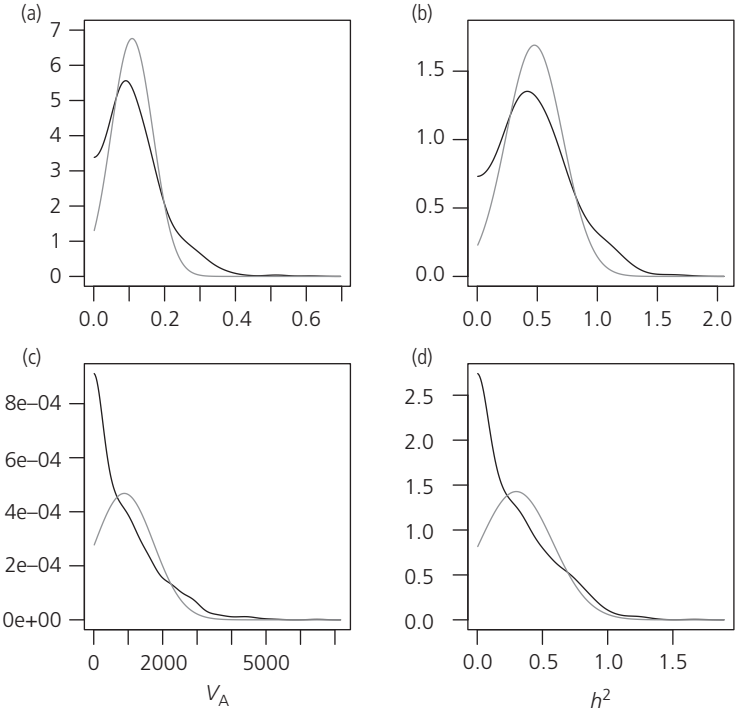

Figure 14.1 Additive genetic variance (a, c) and heritability (b, d) of pronotum length (a, b) and ovary mass (c,d) in field crickets from Simons and Garcia-Gonzalez's (2007) breeding experiment. Black lines show posterior distributions of the parameters, and grey lines show frequentist approximations of the sampling error of their REML estimators, i.e. normal distributions with SD = standard error of each parameter.

Table 14.1 Bayesian and frequentist parameter values from mixed model (sire model) analysis of the cricket dataset. Directly modelled parameters are: the standard deviations associated with sire σ_s, dam σ_d, and the residual variation σ_e. Derived parameters are the additive genetic variance σ^2_{A}, the maternal non-genetic variance σ^2_{M}, the total phenotypic variance σ^2_{P}, and the heritability h^2

Parameter	Posterior mode	Posterior mean	SD of posterior	REML estimate	REML SE
(a) Directly modelled parameters; pronotum length					
σ^2_s	0.024	0.030	0.020	0.027	0.015
σ^2_d	<0.001	0.014	0.013	0.012	0.012
σ^2_e	0.182	0.191	0.016	0.189	0.016
(b) Derived parameters; pronotum length					
σ^2_{A}	0.095	0.121	0.078	0.108	0.059
σ^2_{M}	–0.022	–0.016	0.026	–0.015	0.021
σ^2_{P}	0.231	0.235	0.022	0.228	0.011
h^2	0.43	0.50	0.28	0.48	0.24
(c) Directly modelled parameters; ovary mass					
σ^2_s	2	251	235	222	213
σ^2_d	842	837	268	778	255
σ^2_e	1985	1989	183	1961	170
(d) Derived parameters; ovary mass					
σ^2_{A}	10	1004	942	886	853
σ^2_{M}	626	586	411	557	395
σ^2_{P}	3047	3078	310	2960	340
h^2	<0.01	0.32	0.276	0.30	0.28

SD = standard deviation; REML = restricted maximum likelihood; SE = standard error.

$\sigma^2_e = 0.189$. This yields similar estimates of heritability to the parameter values we obtained from the Bayesian implementation (Table 14.1, Figure 14.1). However, inference of statistical uncertainty in the frequentist parameter estimates is much more challenging (see also Section 14.1); software may calculate *normal approximations of* standard errors (SEs) automatically, but normal approximations to the sampling distribution do not necessarily have sensible interpretations. The basic mechanics require that we obtain an estimate of the sampling variance–covariance matrix of the parameters (i.e. the variances about the estimates of σ^2_s, σ^2_d, and σ^2_e, and the covariances between each estimate). This is obtainable from the model fitted in ASReml-R (with some difficulty; ASReml-R returns the inverse of the average information matrix, which is a key parameter in the algorithm that ASReml uses to solve mixed models, for parameter estimates that are subject to internal scaling. We do not dwell here on the scaling, but we have done it). The sampling variances of the estimates of the variance components are σ^2_{A}: 0.0034, σ^2_{M}: 0.0004, and σ^2_{P}: 0.0001. We can obtain the sampling variance–covariance matrix of the derived parameters σ^2_{A} and σ^2_p by $\boldsymbol{A}\sigma^x\boldsymbol{A}^{\mathrm{T}}$, where σ^x is the variance–covariance matrix of σ^2_s, σ^2_d, and σ^2_e, and $\boldsymbol{A}$ contains the coefficients describing the linear combination of the estimated parameters that is required for derivation of the parameter of interest. Thus, $\boldsymbol{A} = \left(\begin{smallmatrix} 4 & 0 & 0 \\ 1 & 1 & 1 \end{smallmatrix}\right)$, specifying the additive genetic variance as four times the sire variance plus zero times each of the other variances, and specifying the phenotypic variance as the sum of all variance components.

The sampling variance of the estimated heritability is approximately $(\sqrt{\frac{\sigma^2[\sigma_a^2]}{\sigma_p^2}})^2 + (-\sigma_a^2\sqrt{\frac{\sigma^2[\sigma_p^2]}{(\sigma_p^2)^2}})^2 + 2\sigma_{a,p}\sqrt{\frac{\sigma^2[\sigma_a^2]}{\sigma_p^2}}(\frac{-\sigma_a^2\sqrt{\sigma^2[\sigma_p^2]}}{(\sigma_p^2)^2})$, where $\sigma^2[x]$ represents the sampling variance (SE^2) of x. As such, the SE (square root of the sampling variance) of σ_A^2 is 0.059 and the SE of the heritability is 0.24, for pronotum length (Table 14.2).

The tediously derived normal approximations to the sampling errors of the genetic variances and heritabilities of pronotum length and ovary mass are plotted with the posterior distributions of the Bayesian parameter estimates in Figure 14.1. Both approaches to making inferences about parameter values, and to obtaining information about how confident we can be about those inferences, tell us about the same thing. The data available suggest that the heritability of pronotum length is around 0.5, and that the heritability of ovary mass is probably lower. Also, both approaches suggest that our confidence that the real parameter values associated with both traits are close to our best inferences is weak. This is simply a result of the modest sample size.

The consideration of the full posterior distributions of parameters of biological interest that one is more inclined to do after application of a Bayesian analysis is quite useful. First, whilst it must be kept in mind that there are important ways in which a given posterior distribution may be influenced by model structure and prior specification (in addition to the data), a Bayesian posterior distribution is more interpretable as a complete representation of the statistical support for a given parameter having a given value. Where data are scarce, as is often the case for traits of ecological interest in wild populations, consideration of the full posterior distribution can lead to much more statistically sensible interpretations than might otherwise be made. For example, one might conclude that 'the most probable values of V_A and h^2 of cricket ovary mass are very low', but at the same time one should note that 'high values of V_A and h^2 of cricket ovary mass cannot be ruled out'. This latter component of the interpretation is key because there is appreciable density of the posterior distribution at high values of the genetic variances and heritabilities, even if the most probable values are lower.

There are also important differences between the two approaches to assessing uncertainty. First, the Bayesian posterior distribution is exact, given the priors, the data, and the model construction (although the MCMC implementation is an approximation, we expect it to yield the 'true' shape of the distribution). The frequentist approach we took to obtain SEs from the REML analysis is very fundamentally approximate. There is no need to dwell on the specific formulae applied, but the formula for the sampling variance is a very standard approximation, and more importantly, the whole process of obtaining the REML SEs assumes that sampling error of the directly estimated and the subsequently derived parameters, i.e. σ_s^2, σ_d^2, and σ_e^2, and then σ_A^2, σ_P^2 and h^2 are normally distributed. This assumption is of course fundamentally untrue, since normal distributions give non-zero density to all real values, whilst variance components are bounded at zero and heritability is bounded at zero and one. Furthermore, derivation of parameters such as σ_A^2, σ_P^2 and h^2 is very simple, and more complicated parameters that we might be interested in modelling will be totally intractable outside of a Bayesian context.

By roundabout means, we have obtained our first important message about the Bayesian quantitative genetic analysis of natural populations. Transformations of posterior distributions of directly modelled parameters yield valid posterior distributions of derived parameters. This is a powerful and practical feature of Bayesian analysis. However, it can also be dangerous. The details of how a model is specified, including prior specification, can have substantial and even dramatic influence both on directly modelled parameters and on derived parameters (for a discussion of the importance of priors, see Box 14.2).

14.2.3 Alternate model specifications: 'customising' the sire model, and enter the animal model

Inspection of the parameter estimates from the sire model reveals two incongruous results. First, there is non-zero posterior density (i.e. some statistical

Box 14.2 Prior distributions

One of the most problematic issues for new users of Bayesian inference is to deal with priors. This is for two reasons: first, one might think that we are not allowed to have prior knowledge before analysing the data; second, one can struggle to find which prior is the more sensible for one's analysis. Although a full Bayesian analysis includes constructing priors from previous experience, it is not the most popular practice because it requires that previous data exists, and objective rules to construct prior distributions from them.

Common practice is to seek analyses that have only weak prior influence, by seeking *non-informative prior distributions*, which in some cases can include *flat priors* assigning every event the same probability or *diffuse or vague priors* which are also very flat due to low precision (commonly: Gaussian prior with huge variance for fixed effects). The idea of *weakly informative priors* may be a more generally useful concept (Gelman *et al.* 2008). However, regardless of prior specification, it is good practice to check the *sensitivity to the prior distribution*, by running several concurrent priors and checking their influence on the posterior distributions. An example of prior sensitivity analysis applied to heritability is presented in de Villemereuil *et al.* (2013) Appendix B.

It is impossible to provide simple and general advice to beginner Bayesian analysts, but here are a couple of hints. Especially when using BUGS (or JAGS) one would be well advised to read Gelman (2006), as it is especially useful for understanding prior influence on variance components. If conducting quantitative genetic analyses of modest datasets using MCMCglmm (Hadfield 2010), it will be useful to consult the documentation for information about 'parameter-expanded priors' for variance components.

support) for values of heritability that exceed one (Figure 14.1) for both traits. Second, the dam variance is less than the sire variance (Table 14.1). Neither of these makes biological sense (in the latter case, assuming paternal effects are absent or are much smaller than maternal effects), and both occur despite the fact that the sire and dam sources of phenotypic variance have very simple biological interpretations. Being a proportion, heritability cannot exceed one; and the dam variance is the sum of genes and environments of mothers, and neither of these is likely to cause full sibs to resemble one another less than at random (assuming that unmodelled processes such as varying maternal allocation with birth order, or asymmetric sibling competition are not happening). However, the mixed model is blind to the biological interpretations, and since we have specified it in a way that the sire variance can exceed one quarter (which is not biologically meaningful as $V_{\text{sire}} = \frac{1}{4}V_{\text{A}}$, and $V_{\text{A}} \leq V_{\text{P}}$) of the phenotypic variance, and where the dam variance does not have a lower bound at the sire variance, the curious outcomes are not mathematically unsound.

If we have prior belief that we have conducted our analysis appropriately with regard to the interpretation of the sire and dam variances, we can reparameterise the sire model to reflect this belief. The statistical parameters that we directly modelled pertain to the variation associated with dam, sire and individuals. We can redefine these in terms of how the biologically interesting parameters about different sources of variation must be related to one another by

$$\sigma_{\text{P}} \sim U(0, \sigma_{\text{big}}) \tag{14.4a}$$

$$h^2 \sim U(0, 1) \tag{14.4b}$$

$$\omega_d \sim U(0, 1) \tag{14.4c}$$

$$\sigma_s = \sqrt{\frac{h^2\sigma_{\text{P}}^2}{4}} \tag{14.4d}$$

$$\sigma_d = \sqrt{\frac{h^2\sigma_{\text{P}}^2}{4} + \omega_d(1-h^2)\sigma_{\text{P}}^2} \tag{14.4e}$$

$$\sigma_e = \sqrt{\sigma_{\text{P}}^2 - \frac{h^2\sigma_{\text{P}}^2}{2} - \omega_d(1-h^2)\sigma_{\text{P}}^2} \tag{14.4f}$$

For lack of a standardised symbol to denote the proportion of the non-genetic phenotypic variance that is attributable to maternal identity, we denote this parameter ω_d. Expressions 14.5a,b,c simply define the biologically meaningful parameters as having reasonable ranges. The heritability is some fraction of the phenotypic variance, and the maternal

variance is some fraction of the non-genetic variance. Eq. 14.5 simply translates these biologically more sensible parameters into the parameters of the sire model above. This small modification of the sire model to improve biological interpretability is straight-forward given available Bayesian MCMC tools.

Another alternative parameterisation, and ultimately one that is very valuable and general, would be to directly model the process of inheritance of a multivariate trait, rather than modelling dam and sire effects, knowing that they represent some fraction of the genetic variation. This is the approach taken by the 'animal model'. In the animal model, variances associated with parental effects are replaced with a direct model of the effects of an individual's genes, i.e. a model of variation in *breeding values*. A breeding value is defined as twice the expected deviation of an individual's offspring's phenotype from the population mean. More intuitively, it can be thought of as the sum of the effects on an individual's phenotype of the genetic variants carried in that individual's genome. Specifically, an individual's breeding value is

$$a_i \sim N\left(\frac{a_{d_i} + a_{s_i}}{2}, \frac{\sigma^2_A}{2}\right) \quad (14.5)$$

i.e. an individual's expected breeding value is the mean of the breeding values of its parents, and its realized breeding value comes from a distribution with a variance that is half the additive genetic variance in the population. This additional variance is the segregational variance—it can be thought of as reflecting the fact that full siblings are not identical, i.e. full sibs differ in their particular proportional composition of grandparental alleles at different points in their genome, because of the segregation of these alleles into gametes that occurred in their parents. The animal model can still account for maternal effects: we include a dam effect as in the sire model (Eq. 14.1); but this effect is now directly interpretable as the non-(direct) genetic component of the maternal effect, because the genetic component is represented in Eq. 14.5.

Implementing the animal model in BUGS code can be relatively easy:

```
model{
  #priors
  mu~dnorm(0,0.001)
  sigma_a~dunif(0,sigma_big)
  sigma_m~dunif(0,sigma_big)
  sigma_e~dunif(0,sigma_big)

  #random effects
  for(j in 1:N_d){
     a_d[j]~dnorm(0,1/sigma_a^2)
     m_d[j]~dnorm(0,1/sigma_m^2)
   }
   for(k in 1:N_s){
     a_s[k]~dnorm(0,1/sigma_a^2)
   }

   #data
   for(i in 1:N_i){
     a_i[i]~dnorm((a_d[dam[i]]+a_s[sire[i]])/2,1/(0.5*sigma_a^2))
     y[i]~dnorm(mu+a_i[i]+m_d[dam[i]],1/sigma_e^2)
   }
}
```

The key lines are: 10, 11 and 14, where dams and sires are given genetic and non-genetic values sampled from normal distributions with estimated variances on lines 4 and 5; line 19, where offspring breeding values are modelled based on the mid-parent breeding value and the segregational variance following Eq. 14.5; and line 20, where offspring phenotypes are modelled based on the population mean, genetic and maternal effects, and the residual variance.

The posterior distributions of heritability and variance components for cricket pronotum length based on the sire model, the constrained sire model, and the animal model are shown in Figure 14.2. These distributions are all obtained from an identical dataset, with identical sets of assumptions for modelling the relationship between resemblance of relatives and genetic variation, i.e. the infinitesimal model (Chapter 2, Postma; Falconer & Mackay 1996). In analysis of data from single generation, the animal model uses no more information about the genetic basis of variation than does the sire model. This is because variation due to maternal genes, and variation due to maternal identity are perfectly confounded in a single generation study. Thus, a closer look at the differences in the posterior distributions of the parameter estimates in Figure 14.2 provides an opportunity to better understand the models that we have implemented.

Two aspects of the posterior distribution of the heritability of pronotum length differ between the original sire model, and the constrained sire model and the animal model. In the first, there is non-zero posterior density at values greater than one, and

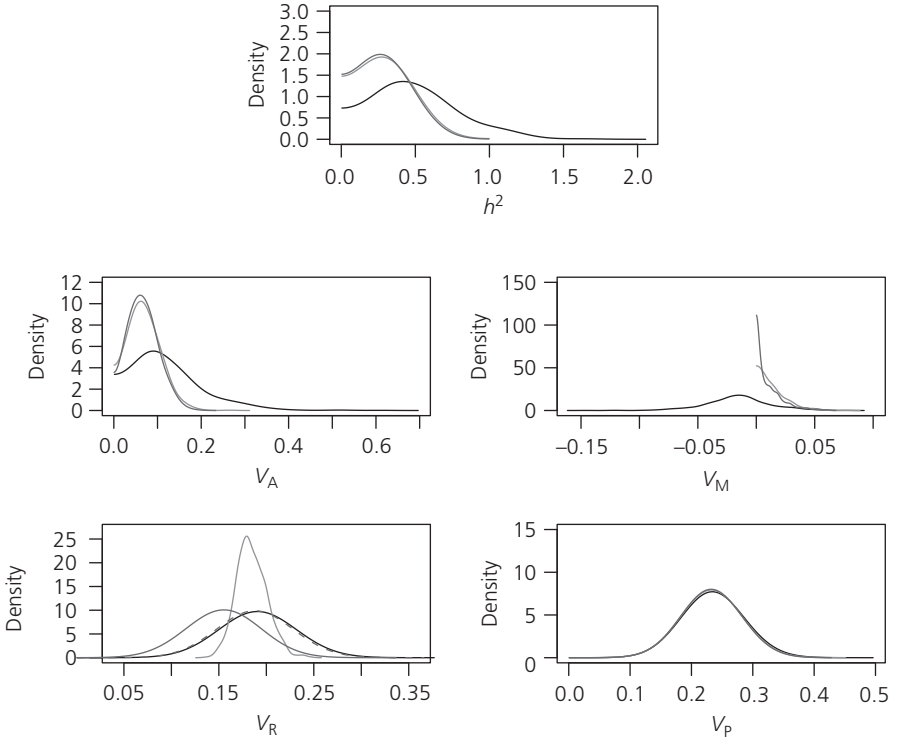

Figure 14.2 Posterior distributions of variance components for cricket pronotum heritability. Black lines represent estimates made with a sire model (as per Eqs. 14.1–3), red lines represent estimates made with a constrained sire model (reparameterisation following Eq. 14.4), and blue lines represent estimates made with an animal model (additive genetic effects directly modelled following Eq. 14.5). The residual variance in a sire model includes segregational variance, whilst it does not in an animal model; the dotted blue line shows the posterior distribution of the sum of the residual and the segregational variance.

the heritability estimates (and the estimates of the genetic variance) are lower in the latter two. The zero density for values of heritability greater than one is due to biologically reasonable structural constraints specifically built into the constrained sire model and inherent in the animal model. However, the lower estimates of σ^2_A and h^2 do not necessarily reflect a shift in the posterior distributions to the left. Rather, because of the different structures of the models, more information is coming into play. Because the reparameterised models explicitly include information about maternal identity in the terms that we relate to genetic variation, and because maternal sibs covary less than paternal sibs (as discussed above, this does not have a simple biological interpretation, but even with the more biologically interpretable models, $\sigma_d < \sigma_s$ remains a property of the data), this lower covariance of maternal sibs results in inference of lower genetic variation.

The final very noticeable difference among the models is that the constrained sire model has a much narrower posterior distribution of the residual variance than the other models. This occurs despite the fact that the constrained sire model and the animal model have similar posterior distributions for all of the other parameters. The reason is simple: in the sire model, genetic variance is derived represented by the sire effect, but only in one quarter proportion—the rest is in the residual; whereas in the re-parameterisations, the genetic variance is fully attributed to the parents, and so in no part represented in the residual. In these relatively simple models, this does not complicate the implementation or interpretation of the models. However, in other analyses, sampling correlations (i.e. non-zero expected error in the estimate of one parameter, given error in another) can become important considerations. We will not dwell on this beyond pointing it out as one of the ways in which simple alternative model parameterisations to analyse the same data can have different properties.

A main take home point from this section is that different model parameterisations can be important for allowing models to be constructed with structures that reflect either classical statistical models, or biological parameters, to a range of different degrees. Sometimes, this is essentially a matter of setting up the priors on the model in different ways. For example, deriving a sire variance from a phenotypic variance, a heritability, and a knowledge that variation due to sires represents one quarter of the additive genetic variance, can be seen as different specifications of the priors: essentially, the constrained sire model reflects the absolute prior belief, i.e. a very strong prior belief, that the heritability cannot exceed one. In other situations, different ways of setting up a model might be viewed more as a different model, rather than as re-parameterisations. In our simple example, analysing a quantitative genetic breeding experiment with a single generation of data, the animal model can be thought of either as a reparameterisation of the constrained sire model, or it can be seen as a different model. In a general pedigree with multiple generations, the animal model would fundamentally make use of more data than a sire model, and so would in fact be a different model beyond prior specification. From a practical perspective in our example analysis, though, the difference is unimportant. The message is the flexibility of Bayesian models, given the availability of tools such as the BUGS language.

14.3 Bayesian quantitative genetic analysis of natural populations: the present and beyond

14.3.1 Quantitative genetic analysis of non-normal quantitative traits

Of the current problems in evolutionary quantitative genetic analysis that will be most fruitfully pursued with the Bayesian toolkit, the analysis of non-normal traits is currently most developed. Traits of evolutionary and ecological significance can take a wide range of statistical distributions, such as binomial, exponential, Poisson (see for example Kruuk *et al.*, this volume) or gamma distributions. Furthermore, trait distributions may be censored, truncated or zero-inflated. The analysis of variation in non-normal traits falls into the Bayesian realm largely by default. Non-Bayesian approaches for fitting generalised mixed models are not yet well developed, as they need to explicitly calculate, or accurately approximate the likelihood, which can

be extremely tricky when for non-normal distributions of response variables (Bolker *et al.*, 2009). The Bayesian framework has the advantage of the MCMC algorithm (see Box 14.1), which is one of the most adaptable estimation algorithm of current statistics. In theory, as long as one can sample from a distribution (even 'exotic' ones like *t*-distribution, or censored or zero-inflated ones), MCMC can be used to sample a Bayesian model. As we saw in the first part of this chapter, the flexibility of Bayesian MCMC-based methods can be a great advantage, leading one to Bayesianism as much due to flexibility as due to philosophy. Here we highlight the use of Bayesian MCMC for quantitative genetic analysis of the types of non-normal data that are often of particular interest in realistic ecological scenarios (e.g. survival, fecundity).

The generalised animal model assumes a hypothetical latent trait θ, which is normally distributed and leads to the observable non-normal trait y through a 'link function' g. The model assumed for θ is identical to a 'classical' animal model

$$\theta = \mu + a + e \tag{14.6a}$$

$$a \sim N(0, \boldsymbol{A}\sigma^2_{\mathrm{A}}) \tag{14.6b}$$

$$e \sim N(0, \mathbf{I}\sigma^2_{\mathrm{R}}) \tag{14.6c}$$

where $\boldsymbol{A}$ is the relatedness matrix (generally derived from a pedigree). Letting y follow any distribution, noted π, which has parameters $g^{-1}(\theta)$ and φ, we can write

$$y \sim \pi(g^{-1}(\theta), \varphi) \tag{14.7}$$

Note that φ is often just a 'nuisance' parameter (for example the data dispersion parameter for a negative binomial distribution), in which we have little interest. Note that with this model, the data have a dispersion linked to the π distribution used, and that σ_{R} is in fact an 'overdispersion' parameter.

For example, for a Poisson distribution (hereafter noted $\mathcal{P}$), we will define Eq. 14.8 as

$$y \sim \mathcal{P}(\log^{-1}(\theta)) \qquad \Longleftrightarrow \qquad y \sim \mathcal{P}(\exp(\theta)) \tag{14.8}$$

using the canonical logarithm link function for the Poisson distribution. Note that no specific dispersion parameter φ (beyond over dispersion, or the variance in e in Eq. 14.6c) is typically applied in the case of the Poisson responses: the variance is equal to the mean for a Poisson distribution.

In addition to statistical convenience, a model such as defined above is very often justifiable, or even desirable, on biological grounds. The latent trait θ represents the additivity of many small sources of variance, including additive genetic effects, as assumed by the infinitesimal model and most of its quantitative genetic applications. Eq. 14.7, however, may be somewhat more arbitrary, in the sense that the link function may be merely a transformation that will be consistent with the distribution of our data, with often little attention to a biological justification. However, sometimes the link function will make good biological sense as well, for example the logarithmic link function for a Poisson trait implies multiplicative effects on the data scale, which may often be natural for skewed count variables.

A further advantage of the model defined above is that it is easily expanded to be multivariate. Indeed, since Eq. 14.6 is defined just as a non-generalised animal model, the multivariate version would be just as easily defined for each latent trait θ_k. Then, each θ_k can be independently transformed into the biological trait y_k via Eq. 14.7. Thus, a multivariate version of this model allows us to estimate the genetic correlation between several traits having different data distribution (for example a binary and a Gaussian trait). For the sake of simplicity and conciseness, we will continue with a univariate model.

When using the generalised animal model to study a quantitative trait, one question arises: on which 'scale' do you want to estimate the heritability? Indeed, regarding Eqs. 14.6 and 14.7, we have three 'traits' on which we can calculate heritability: the data scale y obviously, but also the latent (or linear predictor) scale θ and the link scale $g^{-1}(\theta)$. Each case can be justified: measuring the heritability on the data scale yields the most biologically sensible measure, but with the drawback that it will strongly depend on the distribution of the trait (it does not allow for comparison between Poisson and binomial traits, for example). Because one of the major interests in quantifying heritability lies in its ease for

comparison between traits, it is rarely measured on this scale.

How does one calculate heritability on the different scales? On the data scale, there is actually no easy way to calculate heritability, since we need to have access to the additive genetic variance on this scale, which is not the same as the one on the other scales. The calculation of heritability on the latent scale is however pretty straightforward:

$$h^2_{\theta} = \frac{\sigma^2_{\mathrm{A}}}{\sigma^2_{\mathrm{A}} + \sigma^2_{\mathrm{R}}} \tag{14.9}$$

In order to calculate the heritability on the link scale, we need to take into account the additional variance due to the link function:

$$h^2_{g^{-1}(\theta)} = \frac{\sigma^2_{\mathrm{A}}}{\sigma^2_{\mathrm{A}} + \sigma^2_{\mathrm{R}} + \sigma^2_{\mathrm{link}}} \tag{14.10}$$

Note that, in the case of the above formulated model, the therm σ^2_{R} can be seen as an additive overdispersion parameter (the true residual dispersion being the link-specific variance). Neither scale of calculation of the heritability is either right or wrong. Heritability on the link scale if often sought (e.g. Nakagawa & Schielzeth, 2010; 2013; although the authors use the term 'latent scale' for the here-defined link scale), and may be most natural in many circumstances. To avoid confusion, explicit reporting of formulae used to calculate heritabilities in any particular study will be highly desirable.

We illustrate the study of non-normal traits with the special case of binary traits (such as presence/absence of a character or dichotomous behaviors like allo-/philopatry). We chose binary traits, because they are both a common and a quite difficult type of non-normal trait data, the main issue being the fact that a binary data point conveys little information (because it can only be 0 or 1). Dichotomous phenotypes do not always have a simple Mendelian genetic source, but may be (and most of the time are) the product of a large number of quantitative trait loci. In this case, they fall in the domain of quantitative genetics. The usual model for this kind of traits is the so-called 'threshold model', in which an underlying trait is normally distributed and a threshold is set to separate the dichotomous phenotypes (see Figure 14.3). In a generalised mixed model context, we assume for these traits a normally distributed latent trait θ and a binomial distribution for the actual binary trait:

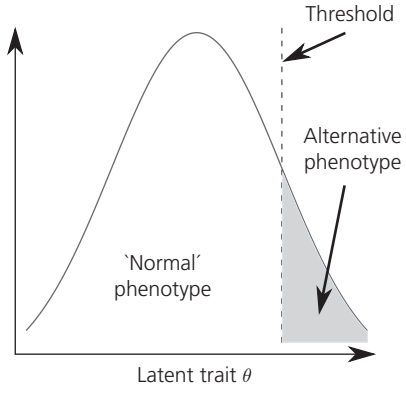

Figure 14.3 A threshold trait is assumed to be produced from a threshold effect depending on the value of the latent trait θ: when θ is below the threshold, the expressed phenotype would be the 'normal' one; when θ is above the threshold, the expressed phenotype would the 'alternative' one.

$$y \sim \mathcal{B}(\mathrm{probit}^{-1}(\theta)) \tag{14.11}$$

where $\mathcal{B}$ is the binomial distribution and the link function used is the probit function[1]. Note that, in this case, the transformation has a straightforward biological meaning. Indeed, when Wright (1934) first introduced the threshold model, he called the latent trait 'liability' with the idea that when this liability was too high (too many deleterious mutations), then the 'alternative' phenotype, which he considered as the 'sick' one, was expressed.

For the sake of comparison between models, the heritability is often estimated on the link scale:

$$h^2 = \frac{\sigma^2_{\mathrm{A}}}{\sigma^2_{\mathrm{A}} + \sigma^2_{\mathrm{R}} + \sigma^2_{\mathrm{link}}} = \frac{\sigma^2_{\mathrm{A}}}{\sigma^2_{\mathrm{A}} + 1 + 1} \tag{14.12}$$

When using binary data, the (overdispersion) residual variance σ^2_{R} in Eq. 14.6 is not identifiable.

[1] Defined as the cumulative density function of a standard normal distribution and, in practice, close to the canonical logit link. The choice of a probit link function is to ensure the continuity with the usual threshold model, which assumes a Gaussian distribution for the 'liability'.

Table 14.2 Comparison of σ^2_A and h^2 estimates for a model with residual variance σ^2_r fixed to 1 or 2. A probit link was used in the model; thus the heritability is calculated as $h^2 = \frac{\sigma^2_A}{\sigma^2_A+\sigma^2_R+1}$

Fixed σ^2_R	σ^2_A	h^2
1	2.00 (0.83)	0.53 (0.080)
2	2.96 (1.07)	0.497 (0.076)

Indeed, binary data cannot be 'overdispersed' in the sense that their variance is solely defined by the proportions p and $1-p$ of each phenotype: $V(y) = p(1-p)$. Thus, the total phenotypic variance is constrained for binary data. Therefore, we can estimate σ^2_A for a given σ^2_R, but not *both* at the same time. Is this a problem? For a unique simulated data set ($h^2 = 0.5$ and 1000 individuals), we compared the MCMC outputs with residual variance fixed to one or two during the estimation process. The results (Table 14.2) show different estimated values for σ^2_A, but consistent results for h^2 (remember that both estimations are on the same data set). Indeed the actual estimated parameter in these models is more the intra-class coefficients which is a class of coefficients the heritability belongs to. Thus the estimate of h^2 is invariant to the arbitrarily chosen fixed value for σ^2_R.

You don't necessarily need Bayesian tools to fit this kind of model. Approximate solutions can be obtained, for example, with the software ASReml (Gilmour *et al.*, 2006). However, when this kind of model is fitted outside of the Bayesian approach, the likelihood is too complex to be easily computed. Thus, two different workarounds are used: i) first consider the binary trait as normally distributed, estimate the heritability using REML, then use a correction (hereafter called corrected REML (REMLc); see Dempster & Lerner, 1950; Lynch & Walsh, 1998); or ii) approximate the likelihood, for example by using penalized quasi-likelihood (PQL; Breslow & Clayton, 1993). de Villemereuil *et al.* (2013) presented a study of the difference in distribution of estimates between these two frequentist and a Bayesian (MCMC) estimation method, based on simulated datasets. Simulations consisted of 1000 replicates of pedigrees and data, for each scenario. Each scenario consisted of one of three true levels of heritability (0.5, 0.3 or 0.1) and two levels of sample size ($n = 1000$ and $n = 200$). As illustrated by Figure 14.4, the PQL estimates are very biased (underestimation of the heritability). This is a known result for binary traits (Goldstein & Rasbash, 1996). The MCMC is a bit more biased than REMLc for small sample size, but is also more precise. The MCMC estimation method is here biased for small sample size because it becomes too sensitive to the particular form of prior used in this study (which was a bit informative; see Box 14.2 on prior sensitivity). Is the high imprecision of the REMLc an issue for estimating heritability? After all, if the standard error and the confidence interval are correctly calculated during the estimation process, we should prefer a non-biased estimator, such as REMLc. However when we calculate the coverage[2] associated with confidence intervals for REMLc and their Bayesian equivalents (credible intervals; CI) for MCMC, we see that MCMC has a better coverage (Figure 14.5, values should lie closely to 95%). Thus, REMLc is altogether imprecise and very confident around its estimation, which is a bad combination.

Bayesian estimation methods (and especially MCMC) are very useful for the study of non-Gaussian traits for two main reasons: i) the MCMC algorithm is flexible enough to allow for almost any kind of distribution to be used as data distribution; and ii), the behaviour of the estimates is always 'correct' according to the model and the prior used (no asymptotic assumptions or approximations have to be made), which means that SEs and associated credible intervals are always a relevant estimation of the imprecision of the estimate[3]. However, as we saw with the binary example, the other side of the coin is that we should be very cautious with the choice of the prior, since for small sample

[2] The coverage is the proportion of time the confidence or credible interval contains the true value of the estimated parameter (here the heritability). The expected coverage corresponds to its nominal value, e.g. 95% for 95% confidence interval.

[3] Again, this is 'according to the prior', which means that a prior too informative might lead to overconfidence in the estimates; but for a Bayesian this is relevant: if you already are quite certain a priori of the results, then you're pretty confident on the estimates you get.

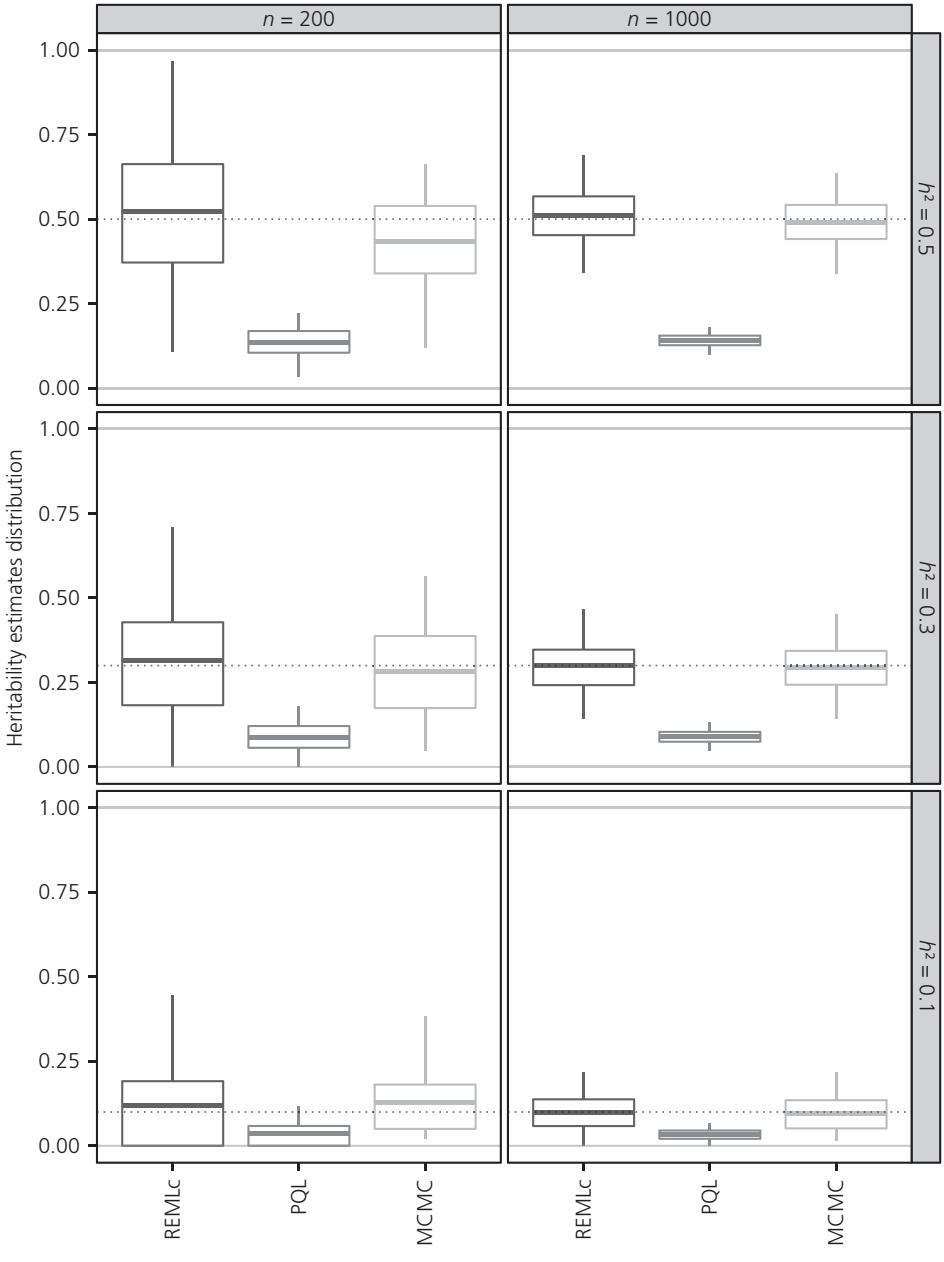

Figure 14.4 Comparison of the heritability estimates distributions for three estimation methods in a simulation study: REMLc, PQL and Bayesian MCMC. Lines show the average value for estimates, boxes show the interquartile interval and whiskers show 95% interquantile interval. Distributions are given for three levels of heritability (0.5, 0.3 and 0.1) and two sample sizes (1000 or 200 individuals).

size, it can have an effect on bias and precision of the estimate (more on this in de Villemereuil *et al.* (2013), Appendix B; Gelman (2006) also provides a useful discussion of the effects of priors for variance components).

14.3.2 Combining quantitative genetic and evolutionary inference

Evolutionary problems, as they play out under real ecological conditions, are inherently complex.

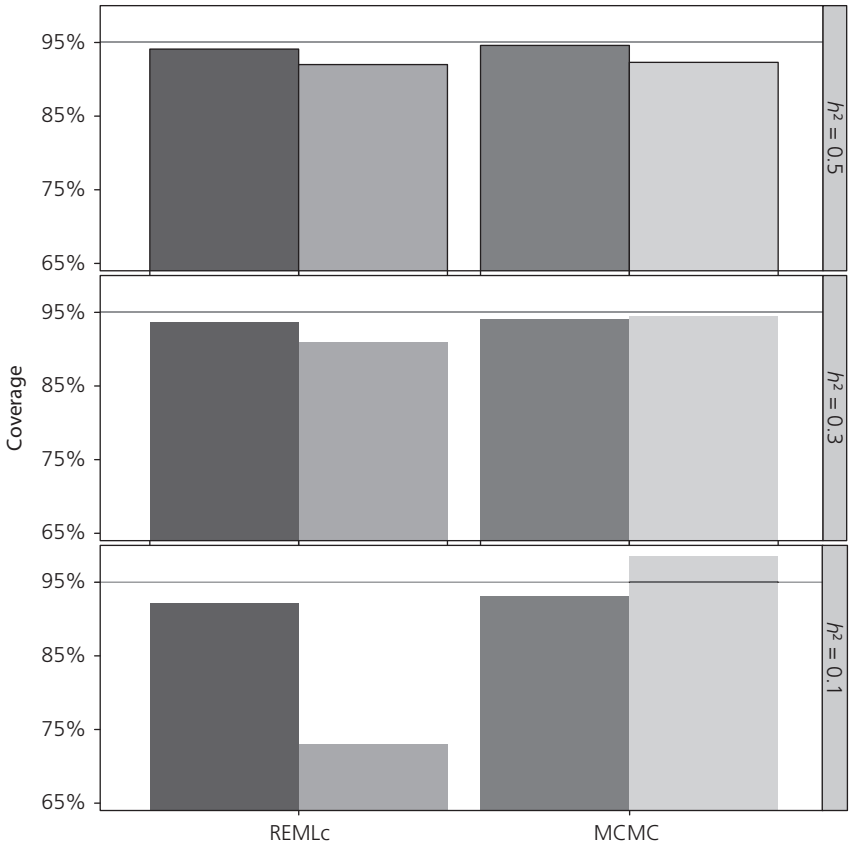

Figure 14.5 Coverage for frequentist REMLc and Bayesian MCMC methods in a simulation study. PQL is not shown, because its coverage is null. Heavy colored bars are for the large sample size (1000 individuals) and light colored bars are for the small sample size (200 individuals).

Consequently, even simplified empirical models will often have to account for more than one process in order to be useful. Bayesian analysis can greatly facilitate the simultaneous evaluation of multiple models, essentially allowing the researcher to break complex problems down into simpler hierarchical levels or 'modules'. In this section, we describe a hierarchical modelling (HM) approach to carry out inference in quantitative genetics. As an example, we illustrate how to calculate the heritability of survival for free–ranging populations in which individuals cannot exhaustively be seen or captured (Lebreton *et al.* 1992; Gimenez *et al.* 2008). Note that survival is intrinsically a non-normal trait and most of the material presented in the previous section applies here. This section is based on work by Papaix *et al.* (2010). In practice, we show how the HM framework allows combining a capture–recapture model to estimate survival whilst accounting for imperfect detection and an animal model to make the decomposition of variance in survival.

Hierarchical modelling (HM) is a powerful approach for analysing complex biological phenomena (Royle & Dorazio 2008; Cressie *et al.* 2009; Buoro *et al.* 2012). One of the most useful applications of HM is when observed data are influenced both by biological processes (e.g. survival probability) and and observation processes (e.g. capture probability). In such a case, a HM can be defined according to three levels: the data at hand Y, the underlying process of interest X and the parameters governing this process. The process X has some distribution governed by a set of parameters θ_X and is generally

not directly or fully observable, e.g. due to issues of detectability or measurement error. Data Y have some distribution that depends on the process X and on a set of parameters θ_Y governing the relationship between Y and X. HM allows modelling the randomness both in the data and in the underlying process via the joint conditional distribution of Y and X given the set of associated parameters θ_Y and θ_X:

$$p(Y, X|\ \theta_X, \theta_Y) = p(Y|\ X, \theta_X) \times p(X|\ \theta_X) \qquad (14.13)$$

where $p(A|B)$ stands for the probability of A given B. This HM formulation is generic and covers a wide variety of models, including so-called state-space models when the process of interest X has a temporal dynamic (e.g. Gimenez *et al.* 2012). HM offers a clear distinction between the biological process and its observation, and so it allows a focus on the former while accommodating uncertainties in the latter.

As an example of HM, let us consider the estimation of survival in capture–recapture models, with the aim of making inferences about the genetic control of variation in survival probability. In free-ranging populations, individual detectability is often less than one. This issue generates data generally collected in the form of 1s and 0s corresponding to a detection or not of I individuals over T sampling occasions. HM has been proposed as a flexible framework to deal with such capture–recapture data (Rivot & Prevost 2002; Gimenez *et al.* 2007; Royle 2008).

In this example, the process X is a binary random variable which represents the demography, with $X_{i,t} = 1$ if individual i is alive and available for detection at time t and 0 if it is dead. If individual i is alive at time $t-1$, it survives until time t with survival probability $\phi_{i,t}$ or dies with a probability $1 - \phi_{i,t}$; in other words,

$$p(X_{i,t}|\ X_{i,t-1}, \phi_{i,t-1}) = Bernoulli(X_{i,t-1} \times \phi_{i,t-1}) \qquad (14.14)$$

Here, survival probability plays the role of θ_X in Eq. 14.13. Now the data Y is a binary random variable, with $Y_{i,t} = 1$ if the individual i is detected at time t and 0 otherwise. These observations are generated from the underlying demographic process, which is partially hidden from the observer, since when an individual is not detected, it is not possible to say whether it is alive or not. If individual i is alive at time t, then it has a probability $p_{i,t}$ of being encountered and a probability $1 - p_{i,t}$ otherwise; in other words, the link between survival and the detection of individuals is made through the observation equation:

$$\left[Y_{i,t}|\ X_{i,t}, p_t\right] = Bernoulli(X_{i,t} \times p_t) \qquad (14.15)$$

Here, the detection probability (p_t) corresponds to the θ_Y in Eq. 14.13.

We have now developed one module: a Bayesian formulation of a simple version of the mark–recapture problem. We can combine this with a second module containing an animal model to bring in quantitative genetic inference based on similarity of relatives. To do so, we follow what was presented in Section 14.3.1. We assume that the random survival process X is related to a continuous underlying latent variable $l_{i,t}$, which, given $X_{i,t-1} = 1$, satisfies:

$$X_{i,t} = \begin{cases} 1 \text{ if } l_{i,t} > \kappa \\ 0 \text{ if } l_{i,t} \leq \kappa \end{cases} \qquad (14.16)$$

for $t = f_i + 1, \ldots, T$, where f_i is the first time individual i is detected, κ is a threshold value, and T is the index of the last interval in time. We assume that the so-called liability $l_{i,t}$ is normally distributed with mean $\mu_{i,t}$ and variance σ_e. To ensure identifiability (because the residual variance of a Bernoulli variable is entirely determined by the mean), and without loss of generality, σ_e is set to 1 and κ to 0.

From this construction (see also Section 14.3.1), we have

$$\phi_{i,t-1} = \Pr(X_{i,t} = 1 | X_{i,t-1} = 1) = F(\mu_{i,t}) \qquad (14.17)$$

where F is the cumulative function of a normal distribution with mean 0 and variance 1. Noting that F^{-1} is the probit function often used to analyse binary data, we can specify an animal model on the mean of the liability:

$$\mu_{i,t} = \text{probit}(\phi_{i,t-1}) = \eta + b_t + e_i + a_i \qquad (14.18)$$

where η is a constant term for the mean survival on the probit scale, b_t is a random yearly effect (i.e. year specific), e_i is an individual random effect which has no genetic basis and a_i is the genetic

value for individual i. Note that the random effect e_i is random among individuals, but not among individuals at different time intervals, because this level of residual variance would be unobservable, i.e. totally confounded with the latent intercept for a Bernoulli response. Covariates can be incorporated as fixed effects possibly affecting survival, e.g. climate effects (Grosbois *et al.* 2008). We assume that the temporal effect b_t is normally distributed with mean zero and variance σ_t^2, e_i normally distributed with mean 0 and variance σ_e^2 whilst the distribution of **a**, the vector of the a_i's, is multivariate normal with mean 0 and variance–covariance matrix $\sigma_A^2 \boldsymbol{A}$, where σ_A^2 is the additive genetic variance and $\boldsymbol{A}$ the additive genetic relationship matrix (see previous sections). Heritability is calculated as the ratio of the additive genetic variance to the total variance:

$$h^2 = \frac{\sigma_A^2}{\sigma_t^2 + \sigma_e^2 + \sigma_A^2 + 1} \qquad (14.19)$$

In order to completely specify the Bayesian model, we provide prior distributions for all parameters. All priors can be selected as sufficiently vague in order to induce little prior knowledge. For the purposes of demonstration, we chose $p \sim U[0,1]$ and $\eta \sim N(0,100)$. We assigned uniform distributions to the SD of the random effects, $\sigma_t \sim U[0,10]$, $\sigma_e \sim U[0,10]$ and $\sigma_A \sim U[0,10]$.

We illustrate this section by estimating the heritability of survival using data from a 29-year study of individually marked blue tits (*Cyanistes caeruleus*) monitored at Pirio, Corsica (see Papaix *et al.* 2010 for more details). The data comprises a total of 614 breeding individuals that were banded, released and recaptured in spring during breeding seasons between 1979 and 2007. The posterior distributions are displayed in Figure 14.6, and the resulting summary estimates are presented in Papaix *et al.* (2010). Detection probability p was high ($p = 0.77$, 95% CI: 0.71–0.82). Survival probability was in agreement with what we were expecting for a small

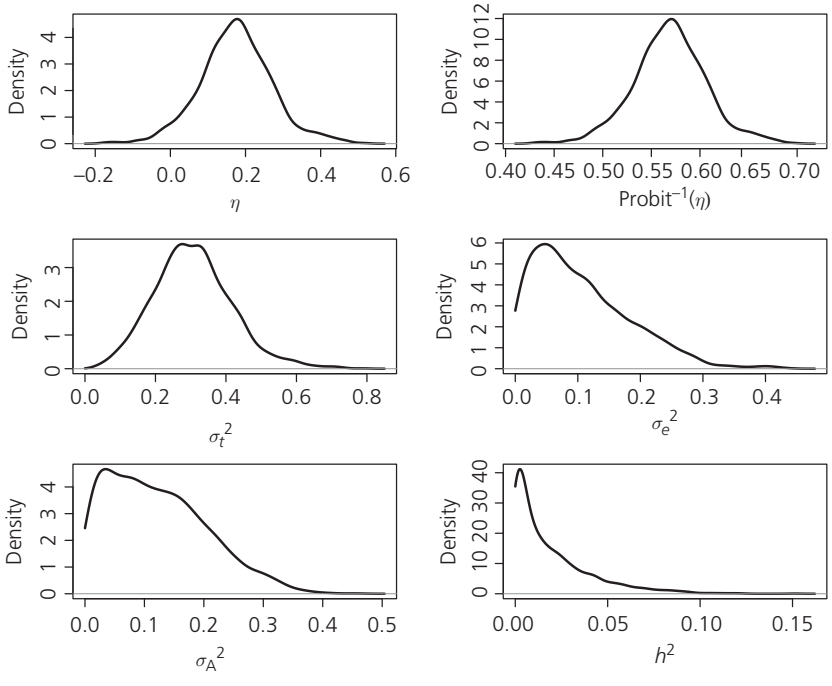

Figure 14.6 Posterior density distributions for parameters of the capture–recapture animal model used for the blue tit data. Notation: η is the mean survival on the probit scale, probit$^{-1}(\eta)$ is the mean survival after back-transformation, σ_t^2 is the variance of the yearly random effect, σ_e^2 is the variance of the non-genetic individual effect, σ_A^2 is the additive genetic variance and h^2 is the heritability.

passerine. The additive genetic variance σ^2_A was low, resulting in a low heritability h^2 (Figure 14.6). The environmental variance σ^2_t was moderate, suggesting temporal variation in survival should not be neglected.

Overall, one can see the HM implementation of this capture–recapture animal model as simple 'modules' being run simultaneously: one module is for the demographic process which is connected to an animal model to decompose variability in survival; the other module is for the observation process and is driven by the detection probabilities. Here, the flow of information is quite intuitive. It is easy to see how longitudinal individual-based data informs the inference of survival probability, and how variation in relatedness provides the basis for genetic inference. It is possible however, and indeed relatively easy, to code arbitrary models where the flow of information through the model structure is not so obvious. Dangerous situations arise easily if model complexity exceeds the intuitivity of the flow of information. Prior information, even if specified in a way that seems uninformative, can easily mask problems in multi-parameter models (Lele 2010, and Lavine 2010, generally); it is possible for innocuous priors to lead to apparent high precision (marked peaks) in posterior distributions of parameters for which there is actually no information. Whilst we generally avoid technical details here, it is worth noting that convergence (the desirable situation where an MCMC algorithm is collecting samples in the model's true region of highest probability density) can be difficult to diagnose as well in complex models. The first line of defence should be biological common sense: follow the flow of information from the data through the model to the parameters.

14.3.3 Towards comprehensive consideration of uncertainty in complex evolutionary analyses

In many empirical studies, multiple statistical procedures are applied to a given dataset. Often the outputs of some procedures serve as inputs for subsequent statistical tests or mathematical procedures. Consideration of statistical uncertainty, e.g. calculation of standard errors, and statistical hypothesis tests are then often conducted assuming that inputs to the very last statistical procedure represent error-free (i.e. statistical sampling error) observations, when in fact they are often themselves statistical parameters or summary statistics that may only be estimated with error. In general 'doing statistics on statistics' will not thoroughly account for statistical uncertainty in any but the last analytical procedures, and so will lead to anticonservative statistical inference. Furthermore, 'doing statistics on statistics', even if the 'statistics upon which statistics are done' are simple, i.e. means from multiple measures at some level of replication, can lead to very severe statistical biases, even when the first steps of statistical analysis seem to be very simple and pragmatic procedures. In this subsection, we present several illustrations of evolutionary quantitative genetic studies (not all exclusively about estimating genetic parameters) where Bayesian approaches have been demonstrated to provide robust analyses in complex problems.

The previous section lays out one potentially comprehensive way of avoiding 'doing statistics on statistics'. Bayesian methods will often facilitate explicit combination of two or more simple models into one more comprehensive model of critically related biological phenomena. More immediately though, samples of posterior distributions of fitted models obtained by MCMC methods have some very convenient properties. They can represent a complete description of (un)certainty of parameter values, and the ways in which uncertainty in one value correlates with uncertainty in another, given the data, the model structure, and the prior specifications. This is a very convenient feature of an applied Bayesian analysis.

A striking recent application in which consideration of the full uncertainty greatly changed the interpretation of a biological result was described by Hadfield *et al.* (2010). Given a fitted animal model, the predicted breeding values can be extracted. Several studies have extracted breeding values from animal model-based analysis of long-term studies, and used them to describe features of the genetics of those populations. A particularly interesting application is to conduct a test of whether or not mean breeding value has changed over time: simply, the regression of breeding values on time provides a test for microevolution.

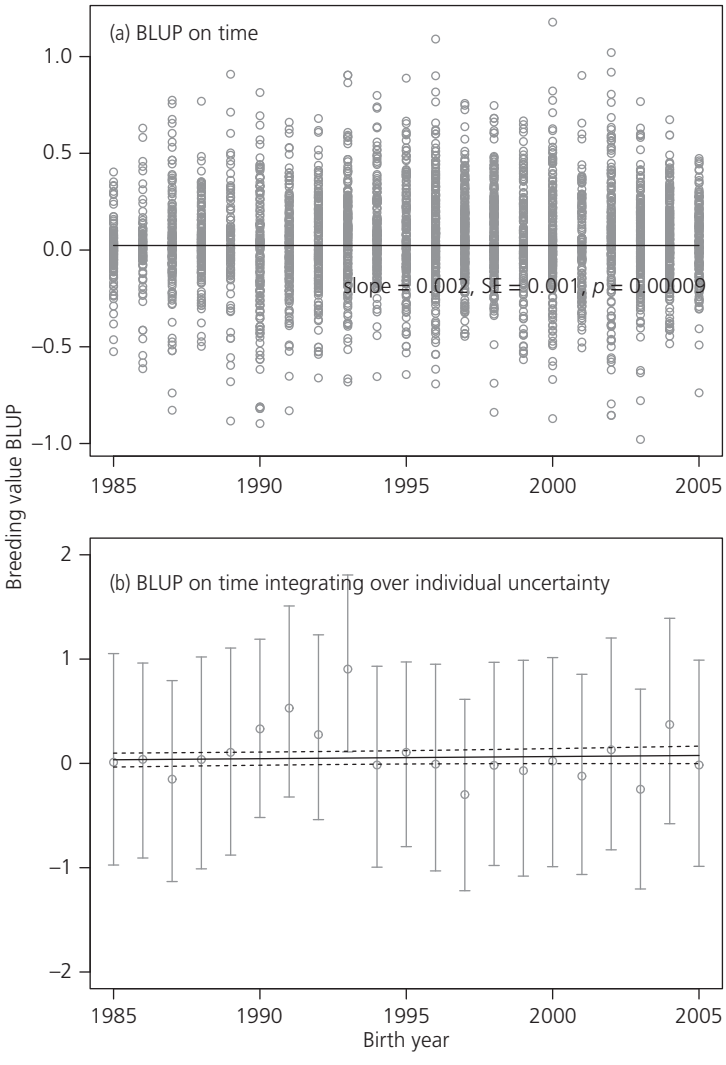

Figure 14.7 Animal model-based inference of microevolution in Soay sheep on St Kilda. (a) shows best linear unbiased predictions (BLUPs) of breeding values of individuals born 1985 to 2005, and their regression on birth year. (b) shows the same regression, but accounting for the fact that each BLUP, i.e. each point in (a), is not a known value, but rather an estimate with uncertainty. Grey points show the BLUP for one randomly chosen individual from each cohort, with 95% CI. The line (with 95% CI of the prediction, integrating over all uncertainty) is the regression corresponding to that in (a); note that this line represents a regression based on all individuals, as in (a), the selection of a single individual per cohort was conducted only for the purpose of plotting. In this example, 18.9% of the posterior distribution of the regression of breeding value on time has a negative slope, corresponding roughly to a two-tailed *P* value of 0.378. Note that Hadfield *et al*. 2010 considered the regression of mean cohort breeding value on time, and we consider the regression of individual breeding value on birth year. The analyses are similar but not identical; our alternative presentation is presented to highlight uncertainty in individual breeding values.

Figure 14.7 shows two regressions. Plot (a) shows the regression of the predicted breeding value of leg length of each Soay sheep (*Ovis aries*) from the ongoing study on the island of Hirta, St Kilda, Outer Hebrides, Scotland (Clutton-Brock & Pemberton 2004), born between 1985 and 2005. In plot (b), the breeding value of a randomly selected individual from each year is plotted with its associated 95% CI for illustration, along with the regression over all individuals in the population (the same line as in plot a). This prediction interval in (b) contains 95% of the density of predictions from regressions of breeding value on year, conducted for each of 1000 samples of the posterior distribution of the same (Bayesian) animal model that was used to get the breeding values in plot (a).

The difference between the two ascertainments of uncertainty in the regression of breeding value on year is very stark. In plot (a), the predicted breeding values are taken to be known values, but as the representative posterior distributions of breeding value in plot (b) show, they are anything but known. Furthermore, individuals that are alive in any given year tend to be closer relatives to other individuals alive at that time, or around that time, than to individuals that lived much earlier or later in the

study. Therefore the very feature of the study that allows genetic parameters to be estimated, the pedigree, causes complex patterns of covariance both in similarity among individuals, but also causes complex patterns of uncertainty in the breeding values of individuals. Integrating over the full posterior distribution of the breeding values in the regression of breeding value on time acknowledges all of the complex patterns of uncertainty in breeding values.

More broadly, there is little tradition in quantitative genetics of considering and reporting uncertainty in estimates of many types of parameters. This is a little bit surprising, given that the field is necessarily so fundamentally statistical. However, it is also a natural result of the predominantly frequentist methods that have dominated the field, and the difficulties that can arise in even approximately describing statistical uncertainty in derived parameters (see Section 14.2.1). Parameters for which SEs, or any other assessment of statistical uncertainty, are rarely provided in quantitative genetic studies include predictions of evolutionary trajectories based on the breeder's or Lande equations (Lande 1979), and descriptions of the geometry of $\boldsymbol{G}$-matrices (reviewed in Walsh & Blows 2009) such as the direction or length of $\mathbf{G}_{\max}$ (the dominant eigenvector of the $\boldsymbol{G}$-matrix), and the constraints that $\boldsymbol{G}$ may impose on adaptive evolution, as potentially described by evolvability and respondability (Hansen & Houle 2008), or metrics of constraint based on the effect of genetic correlations on the rate of adaptation (Stinchcombe and Agrawal 2009). As meta-analysis becomes more important in ecology, genetics, and evolutionary biology in general, there will be an increasing need for estimates of such parameters to be accompanied by metrics of their uncertainty.

Teplitsky *et al.* (2011) inferred the extent to which genetic correlations constrain the rate of adaptation of breeding traits in barn swallows (*Hirundo rustica*), and used Bayesian methods to evaluate the uncertainty in their estimate. One of their goals was to evaluate Agrawal and Stinchcombe 's (2009) metric of constraint due to genetic correlations, which is defined as

$$R = \frac{\Delta \bar{W}_{\mathrm{G}}}{\Delta \bar{W}_{\mathrm{I}}} \tag{14.20}$$

where $\Delta \bar{W}$ is the change in population absolute fitness due to one generation of response to selection, and subscripts G and I denote the change in mean fitness accounting for and discounting genetic correlations, respectively. $\Delta \bar{z}$ is obtained in the standard way according to the Lande (1979) equation $\Delta \bar{z} = \boldsymbol{G}\boldsymbol{\beta}$, and the change in absolute fitness due to evolution is $\Delta W(\Delta \bar{z}) = \Delta \bar{z}^{\mathrm{T}}\boldsymbol{\beta} + \frac{1}{2}\Delta \bar{z}^{\mathrm{T}}\boldsymbol{\gamma}\Delta \bar{\mathbf{z}}$ where $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are vectors and matrices of directional and quadratic selection differentials, respectively.

Clearly, R is a very useful statistic (see also Chapter 12, Teplitsky *et al.*), as it boils down the influences and interactions of multiple aspects of genetic variation, covariation, and their complex relationship with fitness, into a single statistic with a straightforward evolutionary meaning. However, the equations involved in obtaining R represent a rather complex, if biologically highly interpretable, transformation of estimates of genetics and selection, and critically, these estimates are all made with error. Indeed, these estimates, i.e. genetic parameters and aspects of the phenotype–fitness map, are some of the most notoriously difficult parameters to characterise with precision in ecological studies. Teplitsky *et al.* (2011) estimated all of the parameters using Bayesian mixed and ordinary linear models, fitted using the R-package MCMCGLMM (Hadfield 2010), and applied the calculation of R to many effectively independent samples of the posterior solution of mixed models characterising $\boldsymbol{G}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$. The posterior distributions of each component parameter of genetics and selection are shown in Figures 14.8 and 14.9, respectively, and the posterior distribution of R is given in Figure 14.10a.

Morrissey *et al.* (2012) conducted a similar exercise to characterise the effect of genetic correlations on the rate of adaptation of female life-histories in red deer (*Cervus elaphus*). In their analysis, Bayesian techniques allowed demographic theory to be combined with quantitative genetic inferences in order to simultaneously evaluate genetic and selective parameters. As in Teplitsky *et al.* (2011), this allowed inference to be made of the full posterior distribution of R. The inference of R was remarkably similar to that in barn swallows (Figure 14.10), but somewhat less precise. Additionally, Morrissey *et al.* (2012) calculated R under the assumption of the phenotypic gambit, i.e. that phenotypic

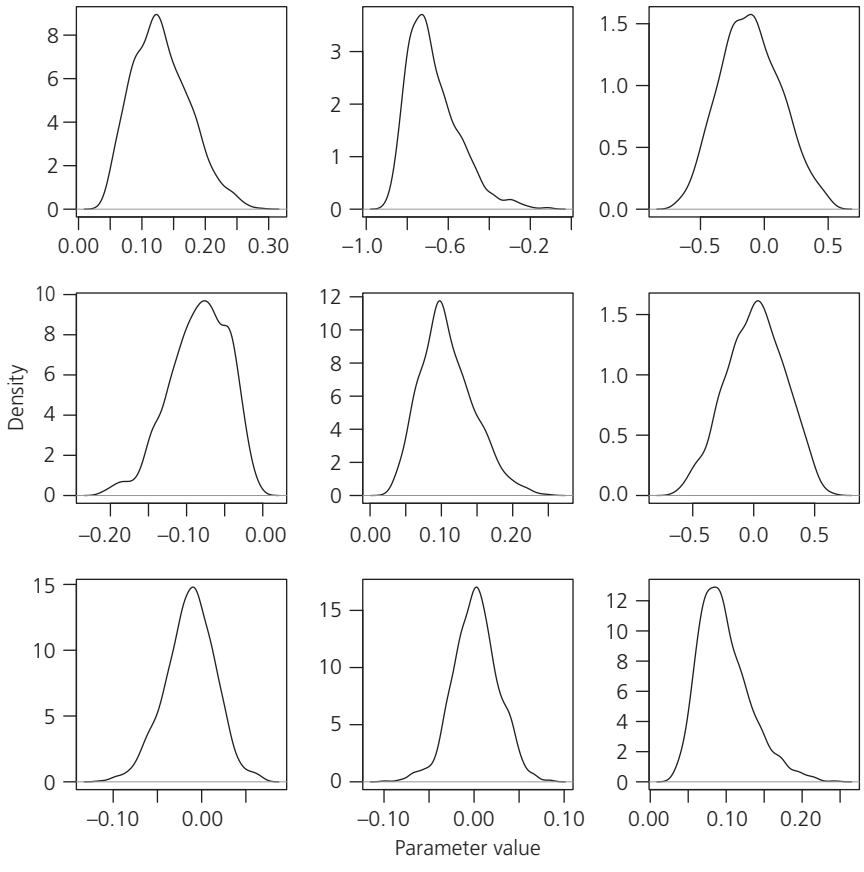

Figure 14.8 Posterior distributions of quantitative genetic parameters of breeding traits in Spanish barn swallows. Traits (top to bottom and left to right) are standardised arrival date, delay before breeding, and clutch size. Diagonal plots are genetic variances, below and above diagonal plots are genetic covariances and correlations, respectively. Posterior distributions are provided by C. Teplitsky from analyses reported in Teplitsky *et al*. (2011).

patterns are substitutable for genetic patterns, and showed that phenotypic correlations do not reveal the influence of genetics on evolutionary trajectories.

The evolutionary constraint metric *R*, and other such parameters, are difficult to characterise with certainty, and also are fundamentally parameters of particular populations at particular times and places, and in particular ecological conditions. Publishing the uncertainty in such metrics will ultimately be necessary to facilitate robust meta-analysis. In turn, meta-analysis will ultimately allow inferences about the general importance of different hypothesised processes, or the mean and range of values of parameters that are difficult to characterise, and that are specific to case studies. In support of meta-analysis, standardised reporting of features of posterior distributions will become highly worthwhile. At the very least, it would be useful if the SDs of the posterior distributions of important parameters were more generally reported (slightly turning a blind eye to philosophy, these may be used as SEs when implementing meta-analyses). Future developments in meta-analysis may allow the complexity of posterior distributions of parameters in individual studies to be accommodated, and so reporting of posterior means, modes, and quantiles (quartiles and 95% ranges may be generally useful) may also eventually prove beneficial.

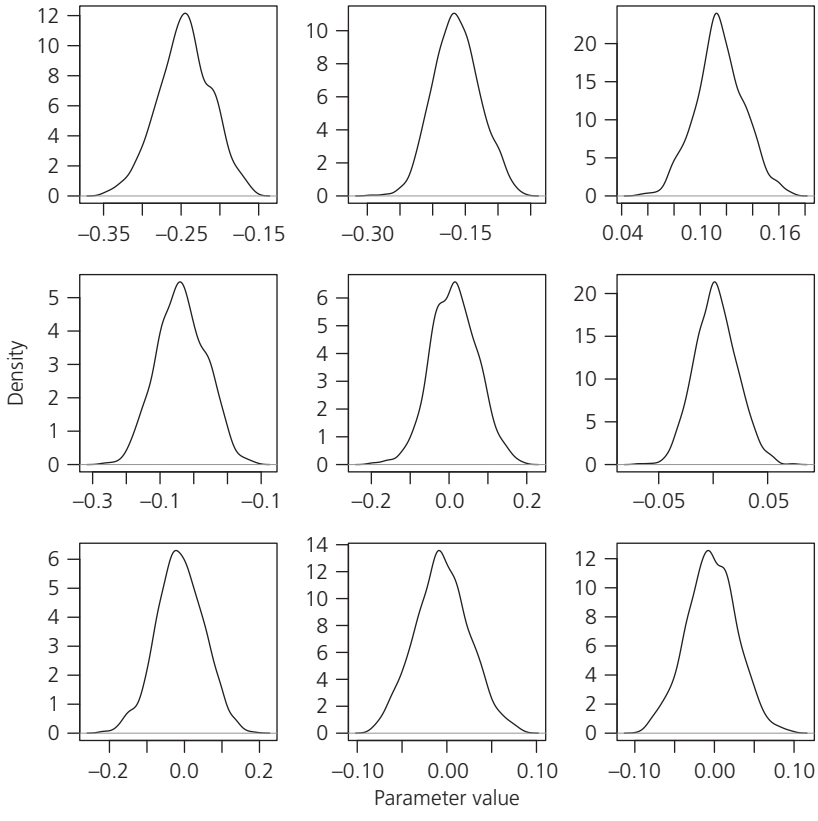

Figure 14.9 Posterior distributions of standardised selection gradients of breeding traits in Spanish barn swallows. Top and middle rows are (left to right) directional and quadratic gradients for arrival date, delay before breeding, and clutch size. Bottom row (left to right) are correlational selection gradients for arrival date and breeding delay, arrival date and clutch size, and breeding delay and clutch size. Posterior distributions are provided by C. Teplitsky from analyses reported in Teplitsky *et al*. (2011).

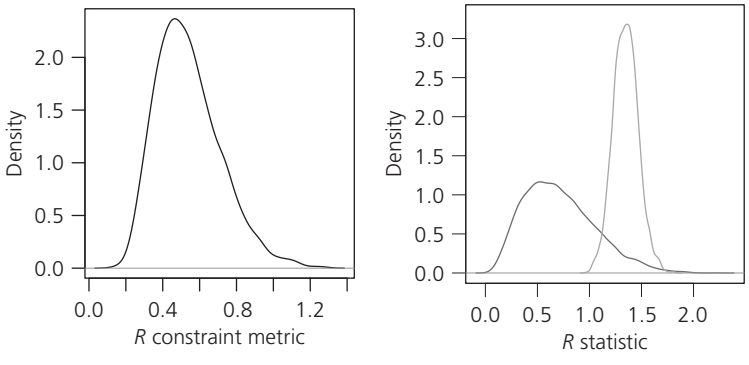

Figure 14.10 The influence of genetic correlations on the rate of adaptation of (a) breeding traits in Spanish barn swallows and (b) female life-history traits in red deer. In (b), the blue function represents the posterior distribution of the constraint metric R, and the red curve represents R calculated under the assumption that $\boldsymbol{P}$ is substitutable for G. The constraint metric R from Agrawal and Stinchcombe (2009) describes the proportion by which genetic correlations change the rate of adaptation (increase of population mean absolute fitness) relative to the rate of adaptation that would occur based on selection gradients and additive genetic variances if genetic correlations were zero. The posterior distribution in (a) is provided by C. Teplitsky from analyses reported in Teplitsky *et al*. (2011).

14.4 Conclusion

We hope that our first section's work through a simple problem highlighted some aspects of the flexibility of Bayesian approaches and currently available associated tools. Our sections on the current applications of Bayesian methodologies are not much more than scattered reports from the frontier of the quantitative genetic analysis of wild populations. For example, Bayesian methods are also contributing greatly to pedigree reconstruction, both in judging uncertainties in pedigrees, and allowing for complex patterns of missingness, whereby sibling relationships can be inferred despite missing the genetic information from the parents. Both of these desirable features of parentage analysis are obtainable without Bayesian methods, but particularly elegant Bayesian solutions have been provided in the software packages MasterBayes (Hadfield *et al.* 2006) and Colony2 (Wang 2004, Jones and Wang 2009). Our theme, throughout both the initial practical demonstration and the subsequent more cutting-edge examples, that Bayesian tools provide pragmatic ways forward in the complex realities of nature, is motivated by a belief that there is much more to come. Other areas where Bayesian methods may provide further benefits for wild quantitative genetics include:

(1) *Incorporation of spatial and temporal structure.* Relatives may generally vary in space for non-genetic reasons, and for reasons that may not be entirely explainable with available data (Stopher *et al.* 2012). Bayesian methods may facilitate much more widespread incorporation of spatial structure in quantitative genetic analyses of data from wild populations.

(2) *Robust modelling.* Data from natural populations tend to contain many outliers. Bayesian methods can generally allow incorporation of very general models of the distribution of any parameters, and this could include the use of thick-tailed distributions where outliers may exist. For example, *t*-distributions may generally be usable where normal distributions are currently more typical. Sorensen and Gianola (2002) discuss robust Bayesian quantitative genetic modelling in some detail.

(3) *More general models of the observation process.* Wild animals are notorious for moving around, and this creates problems for wild quantitative genetics beyond ascertainment of survival (as discussed in Section 14.2). Reproductive success is probably also generally underestimated in nature, and models of the process of observing reproduction (i.e. detectability of mating events and offspring; uncertainty in pedigrees; inference of survival on non-breeders when breeders are most easily censused) could be particularly beneficial. This will be beneficial both for understanding the quantitative genetics of reproduction, and also, considering quantitative genetics more broadly, for ascertainment of fitness in analyses of the selection of quantitative traits in nature.

References

Agrawal, A.F. & Stinchcombe, A.F. (2009) How much do genetic covariances alter the rate of adaptation? *Proceedings of the Royal Society B–Biological Sciences*, **364**, 1593–1605.

Bolker, B.M., Brooks, M.E., Clark, C.J., Geange, S.W., Poulsen, J.R., Stevens, M.H.H. & White, J.-S.S. (2009) Generalized linear mixed models: a practical guide for ecology and evolution. *Trends in Ecology and Evolution*, **24**, 127–135.

Breslow, N.E. & Clayton, D.G. (1993) Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association*, **88**, 9–25.

Brooks, S., Gelman, A., Jones, G.L. & Meng, X.-L. (2011) *Handbook of Markov Chain Monte Carlo*. Chapman and Hall, Boca Raton.

Buoro, M., Prevost, E. & Gimenez, O. (2012) Digging through model complexity: using hierarchical models to uncover evolutionary processes in the wild. *Journal of Evolutionary Biology*, **25**, 2077–2090.

Clutton-Brock, T.H. & Pemberton, J.M., eds. (2004) *Soay sheep: dynamics and selection in an island population*. Cambridge University Press, Cambridge.

Cressie, N., Calder, C., Clark, J.S., Ver Hoef, J.M. & Wikle, C.K. (2009) Accounting for uncertainty in ecological analysis: the strengths and limitations of hierarchical statistical modeling. *Ecological Applications*, **19**, 553–70.

Dempster, E.R. & Lerner, I.M. (1950) Heritability of threshold characters. *Genetics*, **35**, 212–236.

Ellison, A.M. (2004) Bayesian inference in ecology. *Ecology Letters*, **7**, 509–520.

Falconer, D.S. & Mackay, T.F.C. (1989) *Introduction to quantitative genetics*. Longman Group, London.

Gelman, A. (2006) Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, **1**, 515–533.

Gelman, A., Carlin, J.B., Stern, H.S. & Rubin, D.S (2004) *Bayesian data analysis*. Chapman and Hall, New York.

Gelman, A. & Hill, J.S. (2007) *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press, Cambridge.

Gelman, A., Jakulin, A., Pittau, M.A. & Su, Y.S. (2008) A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, **2**, 1360–1383.

Gilmour, A.R., Anderson, R.D. & Rae, A.L. (1985) The analysis of binomial data by a generalized linear mixed model. *Biometrika*, **72**, 593–599.

Gilmour, A.R., Cullis, B.R., Welham, S.J. & Thompson, R. (1999) *ASREML, reference manual*. Biometric bulletin, no 3, NSW Agriculture, Orange Agricultural Institute, Forrest Road, Orange, NSW Australia.

Gilmour, A.R., Gogel, B.J., Cullis, B.R. & Thompson, R. (2006) *ASReml user guide release 2.0*. VSN International Ltd. www.vsni.co.uk/software/asreml/.

Gimenez, O., Lebreton, J.-D., Gaillard, J.-M., Choquet, R. & Pradel, R. (2012) Estimating demographic parameters using hidden process dynamic models. *Theoretical Population Biology*, **82**, 307–316.

Gimenez, O., Rossi, V., Choquet, R., Dehais, C., Doris, B., Varella, H., Vila, J.-P. & Pradel, R. (2007) State-space modelling of data on marked individuals. *Ecological Modelling*, **206**, 431–438.

Gimenez, O., Viallefont, A., Charmantier, A., Pradel, R., Cam, E., Brown, C.R., Covas, R. & Gallard, J.M. (2008) The risk of flawed inference in evolutionary studies when detectability is less than one. *The American Naturalist*, **172**, 441–448.

Goldstein, H. & Rasbash, J. (1996) Improved approximations for multilevel models with binary responses. *Journal of the Royal Statistical Society Series A*, **159**, 505–513.

Grosbois, V., Gimenez, O., Gaillard, J.-M., Pradel, R., Barbraud, C., Clobert, J., Muller, A.P. & Weimerskirch, H. (2008) Assessing the impact of climate variation on survival in vertebrate populations. *Biological Reviews*, **83**, 357–399.

Hadfield, J.D. (2010) MCMC methods for multi-response generalized linear mixed models: the MCMCglmm R package. *Journal of Statistical Software*, **33**, 1–36.

Hadfield, J.D., Richardson, D.S. & Burk, T. (2006) Towards unbiased parentage assignment: combining genetic, behavioural and spatial data in a Bayesian framework. *Molecular Ecology*, **15**, 3715–3731.

Hadfield, J.D., Wilson, A.J., Garant, D., Sheldon, B.C., and Kruuk, L.E.B. (2010) The use and misuse of BLUP in ecology and evolution. *The American Naturalist*, **175**, 116–125.

Hansen, T.F. & Houle, D. (2008) Measuring and comparing evolvability and constraint in multivariate characters. *Journal of Evolutionary Biology*, **21**, 1201–1219.

Jones, O. & Wang, J. (2009) COLONY: a program for parentage and sibship inference from multilocus genotype data. *Molecular Ecology Resources*, **10**, 551–555.

Kruschke, J.K. (2011) *Doing Bayesian data analysis: a tutorial with R and BUGS*. Academic Press New York.

Lande, R. (1978) Quantitative genetic analysis of multivariate evolution applied to brain: body allometry. *Evolution*, **33**, 402–416.

Lavine, M. (2010) Living dangerously with big fancy models. *Ecology*, **91**, 3487.

Lebreton, J.-D., Burnham, K.P., Clobert, J. & Anderson, D.R. (1992) Modeling survival and testing biological hypotheses using marked animals: a unified approach with case studies. *Ecological Monographs*, **62**, 67–118.

Lele, S.R. (2010) Model complexity and information in the data: could it be a house built on sand? *Ecology*, **91**, 3493–3496.

Lunn, D.J., Thomas, A., Best, N. & Spiegelhalter, D. (2000) WinBUGS: a Bayesian modelling framework: concepts, structure and extensibility. *Statistics and Computing*, **10**, 325–337.

Lynch, M. & Walsh, B. (1998) *Genetics and analysis of quantitative traits*. Sinauer, New York.

Morrissey, M.B., Walling, C.A., Wilson, A.J., Pemberton, J.M., Clutton-Brock, T.H. & Kruuk, L.E. (2012) Genetic analysis of life-history constraint and evolution in a wild ungulate population. *American Naturalist*, **179**, E97–114.

Nakagawa, S. & Schielzeth, H. (2010) Repeatability for Gaussian and non Gaussian data: a practical guide for biologists. *Biological Reviews*, **85**, 935–956.

Nakagawa, S. & Schielzeth, H. (2013) A general and simple method for obtaining R2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, **4**, 133–142.

O'Hara, R.B., Cano, J.M., Ovaskainen, O., Teplitsky, C. & Alho, J.S. (2008) Bayesian approaches in evolutionary quantitative genetics. *Journal of Evolutionary Biology*, **21**, 949–957.

Papaix, J., Cubaynes, S., Buoro, M., Charmantier, A., Perret, P. & Gimenez, O. (2010) Combining capture-recapture data and pedigree information to assess heritability of demographic parameters in the wild. *Journal of Evolutionary Biology*, **23**, 2176–2184.

Plummer, M. (2003) JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In:

Proceedings of the 3rd international workship on distributed statistical computing, March 20–22, Vienna, Austria.

R Core Development Team. (2005) *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rivot, E. & Prevost, E. (2002) Hierarchical Bayesian analysis of capture-mark-recapture data. *Canadian Journal of Fisheries and Aquatic Sciences*, **1784**, 1768–1784.

Royle, J.A. (2008) Modeling individual effects in the Cormack-Jolly-Seber Model: a state-space formulation. *Biometrics*, **64**, 364–370.

Royle, J.A. & Dorazio, R.M. (2008) *Hierarchical modeling and inference in ecology: the analysis of data from populations, metapopulations, and communities*. Academic Press, San Diego.

Simmons, L.W. & Garcia-Gonzalez, F. (2007) Female crickets trade offspring viability for fecundity. *Journal of Evolutionary Biology*, **20**, 1617–1623.

Sorensen, D. & Gianola, D. (2002) *Likelihood, Bayesian and MCMC methods in quantitative genetics*. Springer, New York.

Stopher, K.V., Walling, C.A., Morris, A., Guinness, F.E., Clutton-Brock, T.H., Pemberton, J.M. & Nussey, D.H. (2012) Shared spatial effects on quantitative genetic parameters: accounting for spatial autocorrelation and home range overlap reduces estimates of heritability in wild red deer. *Evolution*, **66**, 2411–2426.

Teplitsky, C., Mouawad, N.G., Balbontin, J., de Lope, F. & Moller, A.P. (2011) Quantitative genetics of migration syndromes: a study of two barn swallow populations. *Journal of Evolutionary Biology*, **24**, 2025–2039.

de Villemereuil, P., Gimenez, O. & Doligez, B. (2013) Comparing parent–offspring regression with frequentist and Bayesian animal models to estimate heritability in wild populations: a simulation study for Gaussian and binary traits. *Methods in Ecology and Evolution*, **4**, 260–275.

Walsh, B. & Blows, M.W. (2009) Abundant genetic variation + strong selection = multivariate genetic constaints: a geometric view of adaptation. *Annual Review of Ecology, Evolution and Systemetics*, **40**, 41–59.

Wang, J. (2004) Sibship reconstruction from genetic data with typing errors. *Genetics*, **166**, 1963–1979.

Wright, S. (1934) An analysis of variability in number of digits in an inbred strain of guinea pigs. *Genetics*, **19**, 506–536.